Noname manuscript No.
(will be inserted by the editor)

The Advice Complexity of a Class of Hard Online Problems

Joan Boyar · Lene M. Favrholdt · Christian Kudahl · Jesper W. Mikkelsen

May 26, 2016

arXiv:1408.7033v3 [cs.DS] 25 May 2016

Abstract The advice complexity of an online problem is a measure of how much knowledge of the future an online algorithm needs in order to achieve a certain competitive ratio. Using advice complexity, we define the first online complexity class, AOC. The class includes independent set, vertex cover, dominating set, and several others as complete problems. AOC-complete problems are hard, since a single wrong answer by the online algorithm can have devastating consequences. For each of these problems, we show that $\log\left(1+(c-1)^{c-1}/c^c\right)n = \Theta(n/c)$ bits of advice are necessary and sufficient (up to an additive term of $O(\log n)$) to achieve a competitive ratio of c.

The results are obtained by introducing a new string guessing problem related to those of Emek et al. (TCS 2011) and Böckenhauer et al. (TCS 2014). It turns out that this gives a powerful but easy-to-use method for providing both upper and lower bounds on the advice complexity of an entire class of online problems, the AOC-complete problems.

Previous results of Halldórsson et al. (TCS 2002) on online independent set, in a related model, imply that the advice complexity of the problem is $\Theta(n/c)$. Our results improve on this by providing an exact formula for the higher-order term. For online disjoint path allocation, Böckenhauer et al. (ISAAC 2009) gave a lower bound of $\Omega(n/c)$ and an upper bound of $O((n\log c)/c)$ on the advice complexity. We improve

A preliminary version of this paper appeared in the proceedings of the 32nd International Symposium on Theoretical Aspects of Computer Science (STACS 2015), *Leibniz International Proceedings in Informatics* 30: 116-129, 2015.

This work was partially supported by the Villum Foundation and the Danish Council for Independent Research, Natural Sciences.

J. Boyar · L.M. Favrholdt · C. Kudahl · J.W. Mikkelsen
Department of Mathematics and Computer Science, University of Southern Denmark, 5230 Odense M, Denmark
Tel.: +45 6550-2338,
E-mail: {joan,lenem,jesperwm,kudahl}@imada.sdu.dk

on the upper bound by a factor of $\log c$. For the remaining problems, no bounds on their advice complexity were previously known.

Keywords online algorithms, advice complexity, complexity class, asymmetric string guessing, covering designs, Asymmetric Online Covering (AOC)

1 Introduction

An online problem is an optimization problem in which the input is divided into small pieces, usually called requests, arriving sequentially. An online algorithm must serve each request without any knowledge of future requests, and the decisions made by the online algorithm are irrevocable. The goal is to minimize or maximize some objective function.

Traditionally, the quality of an online algorithm is measured by the competitive ratio, which is an analog of the approximation ratio for approximation algorithms: The solution produced by the online algorithm is compared to the solution produced by an optimal offline algorithm, OPT, which knows the entire request sequence in advance, and only the worst case is considered.

For some online problems, it is impossible to achieve a good competitive ratio. As an example, consider the classical problem of finding a maximum independent set in a graph. Suppose that, at some point, an online algorithm decides to include a vertex v in its solution. It then turns out that all forthcoming vertices in the graph are connected to v, but not to each other. Thus, the online algorithm cannot include any of these vertices. On the other hand, OPT knows the entire graph, and so it rejects v and instead takes all forthcoming vertices. In fact, one can easily show that, even if we allow randomization, no online algorithm for this problem can obtain a competitive ratio better than $\Omega(n)$, where n is the number of vertices in the graph.

A natural question for online problems, which is not answered by competitive analysis, is the following: *Is there some small amount of information such that, if the online algorithm knew this, then it would be possible to achieve a significantly better competitive ratio?* Our main result is a negative answer to this question for an entire class of hard online problems, including independent set. We prove our main result in the recently introduced advice complexity model. In this model, the online algorithm is provided with b bits of advice about the input. No restrictions are placed on the advice. This means that the advice could potentially encode some knowledge which we would never expect to be in possession of in practice, or the advice could be impossible to compute in any reasonable amount of time. Lower bounds obtained in the advice complexity model are therefore very robust, since they do not rely on any assumptions about the advice. If we know that b bits of advice are necessary to be c-competitive, then we know that *any* piece of information which can be encoded using less than b bits will not allow an online algorithm to be c-competitive.

In this paper, we use advice complexity to introduce the first complexity class for online problems. The complete problems for this class, one of which is independent set, are very hard in the online setting. We essentially show that for the complete problems in the class, a c-competitive online algorithm needs as much advice as is required to explicitly encode a solution of the desired quality. One important feature

of our framework is that we introduce an abstract online problem which is complete for the class and well-suited to use as the starting point for reductions. This makes it easy to prove that a large number of online problems are complete for the class and thereby obtain tight bounds on their advice complexity.

1.1 Advice Complexity

Advice complexity [7, 14, 15, 22] is a quantitative and standardized, i.e., problem independent, way of relaxing the online constraint by providing the algorithm with partial knowledge of the future. The main idea of advice complexity is to provide an online algorithm, ALG, with some advice bits. These bits are provided by a trusted oracle, 0, which has unlimited computational power and knows the entire request sequence.

In the first model proposed [14], the advice bits were given as answers (of varying lengths) to questions posed by ALG. One difficulty with this model is that using at most 1 bit, three different options can be encoded (giving no bits, a 0, or a 1). This problem was addressed by the model proposed in [15], where the oracle is required to send a fixed number of advice bits per request. However, for the problems we consider, one bit per request is enough to guarantee an optimal solution, and so this model is not applicable. Instead, we will use the "advice-on-tape" model [7], which allows for a sublinear number of advice bits while avoiding the problem of encoding information in the length of each answer. Before the first request arrives, the oracle prepares an *advice tape*, an infinite binary string. The algorithm ALG may, at any point, read some bits from the advice tape. The *advice complexity* of ALG is the maximum number of bits read by ALG for any input sequence of at most a given length.

When advice complexity is combined with competitive analysis, the central question is: How many bits of advice are necessary and sufficient to achieve a given competitive ratio c?

Definition 1 (Competitive ratio [23, 32] and advice complexity [7, 22]) The input to an online problem, P, is a request sequence $\sigma = \langle r_1, \ldots, r_n \rangle$. An *online algorithm with advice*, ALG, computes the output $y = \langle y_1, \ldots, y_n \rangle$, under the constraint that y_i is computed from $\varphi, r_1, \ldots, r_i$, where φ is the content of the advice tape. Each possible output for P is associated with a *score*. For a request sequence σ, $\text{ALG}(\sigma)$ ($\text{OPT}(\sigma)$) denotes the score of the output computed by ALG (OPT) when serving σ.

If P is a maximization problem, then ALG is $c(n)$-*competitive* if there exists a constant, α, such that, for all $n \in \mathbb{N}$,

$$\text{OPT}(\sigma) \leq c(n) \cdot \text{ALG}(\sigma) + \alpha,$$

for all request sequences, σ, of length at most n. If P is a minimization problem, then ALG is $c(n)$-*competitive* if there exists a constant, α, such that, for all $n \in \mathbb{N}$,

$$\text{ALG}(\sigma) \leq c(n) \cdot \text{OPT}(\sigma) + \alpha,$$

for all request sequences, σ, of length at most n. In both cases, if the inequality holds with $\alpha = 0$, we say that ALG is *strictly $c(n)$-competitive*.

The *advice complexity, $b(n)$, of an algorithm*, ALG, is the largest number of bits of φ read by ALG over all possible inputs of length at most n. The *advice complexity of a problem*, P, is a function, $f(n,c)$, $c \geq 1$, such that the smallest possible advice complexity of a strictly c-competitive online algorithm for P is $f(n,c)$.

In this paper, we only consider deterministic online algorithms (with advice). Note that both the advice read and the competitive ratio may depend on n, but, for ease of notation, we often write b and c instead of $b(n)$ and $c(n)$. Also, by this definition, $c \geq 1$, for both minimization and maximization problems. For minimization problems, the score is also called the *cost*, and for maximization problems, the score is also called the *profit*. Furthermore, we use output and *solution* interchangeably. Lower and upper bounds on the advice complexity have been obtained for many problems, see e.g. [2, 4–10, 13–15, 17, 18, 22, 24, 26, 28, 30, 31].

1.2 String guessing

In [5, 15], the advice complexity of the following *string guessing* problem, SG, is studied: For each request, which is simply empty and contains no information, the algorithm tries to guess a single bit (or more generally, a character from some finite alphabet). The correct answer is either revealed as soon as the algorithm has made its guess (*known history*), or all of the correct answers are revealed together at the very end of the request sequence (*unknown history*). The goal is to guess correctly as many bits as possible.

The problem was first introduced (under the name generalized matching pennies) in [15], where a lower bound for randomized algorithms with advice was given. In [5], the lower bound was improved for the case of deterministic algorithms. In fact, the lower bound given in [5] is tight up to lower-order additive terms. While SG is rather uninteresting in the view of traditional competitive analysis, it is very useful in an advice complexity setting. Indeed, it has been shown that the string guessing problem can be reduced to many classical online problems, thereby giving lower bounds on the advice complexity for these problems. This includes bin packing [10], the k-server problem [18], list update [9], metrical task system [15], set cover [5] and a certain version of maximum clique [5].

1.2.1 Asymmetric string guessing

In this paper, we introduce a new string guessing problem called *asymmetric string guessing,* ASG, formally defined in Section 2. The rules are similar to those of the original string guessing problem with an alphabet of size two, but the score function is asymmetric: If the algorithm answers 1 and the correct answer is 0, then this counts as a single wrong answer (as in the original problem). On the other hand, if the algorithm answers 0 and the correct answer is 1, the solution is deemed infeasible and the algorithm gets an infinite penalty. This asymmetry in the score function forces the algorithm to be very cautious when making its guesses.

As with the original string guessing problem, ASG is not very interesting in the traditional framework of competitive analysis. However, it turns out that ASG captures, in a very precise way, the hardness of problems such as online independent set and online vertex cover.

1.3 Problems

Many of the problems that we consider are graph problems, and most of them are studied in the *vertex-arrival model*. In this model, the vertices of an unknown graph are revealed one by one. That is, in each round, a vertex is revealed together with all edges connecting it to previously revealed vertices. For the problems we study in the vertex-arrival model, whenever a vertex, v, is revealed, an online algorithm ALG must (irrevocably) decide if v should be included in its solution or not. Denote by V_{ALG} the vertices included by ALG in its solution after all vertices of the input graph have been revealed. The individual graph problems are defined by specifying the set of feasible solutions. The cost (profit) of an infeasible solution is ∞ ($-\infty$).

The problems we consider in the vertex-arrival model are:

- ONLINE VERTEX COVER. A solution is feasible if it is a vertex cover in the input graph. The problem is a minimization problem.
- ONLINE CYCLE FINDING. A solution is feasible if the subgraph induced by the vertices in the solution contains a cycle. We assume that the presented graph always contains a cycle. The problem is a minimization problem
- ONLINE DOMINATING SET. A solution is feasible if it is a dominating set in the input graph. The problem is a minimization problem.
- ONLINE INDEPENDENT SET. A solution is feasible if it is an independent set in the input graph. The problem is a maximization problem.

We emphasize that the classical 2-approximation algorithm for offline vertex cover cannot be used in our online setting, even though the algorithm is greedy. That algorithm greedily covers the edges (by selecting both endpoints) one by one, but this is not possible in the vertex-arrival model.

Apart from the graph problems in the vertex-arrival model mentioned above, we also consider the following online problems. Again, the cost (profit) of an infeasible solution is ∞ ($-\infty$).

- ONLINE DISJOINT PATH ALLOCATION. A path with $L+1$ vertices $\{v_0,\ldots,v_L\}$ is given. Each request (v_i,v_j) is a subpath specified by the two endpoints v_i and v_j. A request (v_i,v_j) must immediately be either accepted or rejected. This decision is irrevocable. A solution is feasible if the subpaths that have been accepted do not share any edges. The profit of a feasible solution is the number of accepted paths. The problem is a maximization problem.
- ONLINE SET COVER (set-arrival version). A finite set U known as the *universe* is given. The input is a sequence of n finite subsets of U, $(A_1,\ldots,A_n)$, such that $\cup_{1\leq i\leq n}A_i = U$. A subset can be either accepted or rejected. Denote by S the set of indices of the subsets accepted in some solution. The solution is feasible if

$\cup_{i \in S} A_i = U$. The cost of a feasible solution is the number of accepted subsets. The problem is a minimization problem.

1.4 Preliminaries

Throughout the paper, we let n denote the number of requests in the input.

We let log denote the binary logarithm $\log_2$ and ln the natural logarithm $\log_e$.

By a *string* we always mean a bit string. For a string $x \in \{0,1\}^n$, we denote by $|x|_1$ the Hamming weight of x (that is, the number of 1s in x) and we define $|x|_0 = n - |x|_1$. Also, we denote the i'th bit of x by x_i, so that $x = x_1 x_2 \ldots x_n$.

For $n \in \mathbb{N}$, define $[n] = \{1, 2, \ldots, n\}$. For a subset $Y \subseteq [n]$, the *characteristic vector* of Y is the string $y = y_1 \ldots y_n \in \{0,1\}^n$ such that, for all $i \in [n]$, $y_i = 1$ if and only if $i \in Y$. For $x, y \in \{0,1\}^n$, we write $x \sqsubseteq y$ if $x_i = 1 \Rightarrow y_i = 1$ for all $1 \leq i \leq n$.

If the oracle needs to communicate some integer m to the algorithm, and if the algorithm does not know of any upper bound on m, the oracle needs to use a self-delimiting encoding. For instance, the oracle can write $\lceil \log(m+1) \rceil$ in unary (a string of 1's followed by a 0) before writing m itself in binary. In total, this encoding uses $2\lceil \log(m+1) \rceil + 1 = O(\log m)$ bits. Slightly more efficient encodings exist, see e.g. [6].

1.5 Our contribution

In Section 3, we give lower and upper bounds on the advice complexity of the new asymmetric string guessing problem, ASG. The bounds are tight up to an additive term of $O(\log n)$. Both upper and lower bounds hold for the competitive ratio as well as the *strict* competitive ratio.

More precisely, if b is the number of advice bits necessary and sufficient to achieve a (strict) competitive ratio $c > 1$, then we show that

$$b = \log\left(1 + \frac{(c-1)^{c-1}}{c^c}\right) n \pm \Theta(\log n), \tag{1}$$

where

$$\frac{1}{e \ln 2} \frac{n}{c} \leq \log\left(1 + \frac{(c-1)^{c-1}}{c^c}\right) n \leq \frac{n}{c}.$$

This holds for all variants of the asymmetric string guessing problem (minimization/maximization and known/unknown history). See Figure 1 on page 14 for a graphical plot. For the lower bound, the constant hidden in $\Theta(\log n)$ depends on the additive constant α of the c-competitive algorithm. We only consider $c > 1$, since in order to be strictly 1-competitive, an algorithm needs to correctly guess every single bit. It is easy to show that this requires n bits of advice (see e.g. [5]). By Remark 1 in section 3, this also gives a lower bound for being 1-competitive.

In Section 4, we introduce a class, AOC, of online problems. The class AOC essentially consists of those problems which can be reduced to ASG. In particular, for any problem in AOC, our upper bound on the advice complexity for ASG applies.

This is one of the few known examples of a general technique for constructing online algorithms with advice, which works for an entire class of problems.

On the hardness side, we show that several online problems, including ONLINE VERTEX COVER, ONLINE CYCLE FINDING, ONLINE DOMINATING SET, ONLINE INDEPENDENT SET, ONLINE SET COVER and ONLINE DISJOINT PATH ALLOCATION are AOC-complete, that is, they have the same advice complexity as ASG. We prove this by providing reductions from ASG to each of these problems. The reductions preserve the competitive ratio and only increase the number of advice bits by an additive term of $O(\log n)$. Thus, we obtain bounds on the advice complexity of each of these problems which are essentially tight. Finally, we give a few examples of problems which belong to AOC, but are provably not AOC-complete. This first complexity class with its many complete problems could be the beginning of a complexity theory for online algorithms.

As a key step in obtaining our results, we establish a connection between the advice complexity of ASG and the size of covering designs (a well-studied object from the field of combinatorial designs).

1.5.1 Discussion of results

Note that the offline versions of the AOC-complete problems have very different properties. Finding the shortest cycle in a graph can be done in polynomial time. There is a greedy 2-approximation algorithm for finding a minimum vertex cover. No $o(\log n)$-approximation algorithm exists for finding a minimum set cover (or a minimum dominating set), unless $\mathsf{P} = \mathsf{NP}$ [29]. For any $\varepsilon > 0$, no $n^{1-\varepsilon}$-approximation algorithm exists for finding a maximum independent set, unless $\mathsf{ZPP} = \mathsf{NP}$ [21]. Yet these AOC-complete problems all have essentially the same high advice complexity. Remarkably, the algorithm presented in this paper for problems in AOC is *oblivious* to the input: it ignores the input and uses only the advice to compute the output. Our lower bound proves that for AOC-complete problems, this oblivious algorithm is optimal. This shows that for AOC-complete problems, an adversary can reveal the input in such a way that an online algorithm simply cannot deduce any useful information from the previously revealed requests when it has to answer the current request. Thus, even though the AOC-complete problems are very different in the offline setting with respect to approximation, in the online setting, they become equally hard since an adversary can prevent an online algorithm from using any non-trivial structure of these problems.

Finally, we remark that the bounds (1) are under the assumption that the number of 1s in the input string (that is, the size of the optimal solution) is chosen adversarially. In fact, if t denotes the number of 1s in the input string, we give tight lower and upper bounds on the advice complexity as a function of both n, c, and t. We then obtain (1) by calculating the value of t which maximizes the advice needed (it turns out that this value is somewhere between $n/(ec)$ and $n/(2c)$). If t is smaller or larger than this value, then our algorithm will use less advice than stated in (1).

1.5.2 Comparison with previous results

The original string guessing problem, SG, can be viewed as a maximization problem, the goal being to correctly guess as many of the n bits as possible. Clearly, OPT always obtains a profit of n. With a single bit of advice, an algorithm can achieve a strict competitive ratio of 2: The advice bit simply indicates whether the algorithm should always guess 0 or always guess 1. This is in stark contrast to ASG, where linear advice is needed to achieve any constant competitive ratio. On the other hand, for both SG and ASG, achieving a constant competitive ratio $c < 2$ requires linear advice. However, the exact amount of advice required to achieve such a competitive ratio is larger for ASG than for SG. See Figure 1 for a graphical comparison.

The problems ONLINE INDEPENDENT SET and ONLINE DISJOINT PATH ALLOCATION, which we show to be AOC-complete, have previously been studied in the context of advice complexity or similar models. We present a detailed comparison of our work to these previous results.

In [7], among other problems, the advice complexity of ONLINE DISJOINT PATH ALLOCATION is considered. It is shown that a strictly c-competitive algorithm must read at least $\frac{n+2}{2c} - 2$ bits of advice. Comparing with our results, we see that this lower bound is asymptotically tight. On the other hand, the authors show that for any $c \geq 2$, there exists a strictly c-competitive online algorithm reading at most b bits of advice, where

$$b = \min\left\{ n\log\left(\frac{c}{(c-1)^{(c-1)/c}}\right), \frac{n\log n}{c} \right\} + 3\log n + O(1)\,.$$

We remark that $n\log\left(c/(c-1)^{(c-1)/c}\right) \geq (n\log c)/c$, for $c \geq 2$. Thus, this upper bound is a factor of $2\log c$ away from the lower bound.

In [19], the problem ONLINE INDEPENDENT SET is studied in a multi-solution model. In this model, an online algorithm is allowed to maintain multiple solutions. The algorithm knows (a priori) the number n of vertices in the input graph. The model is parameterized by a function $r(n)$. Whenever a vertex v is revealed, the algorithm can include v in at most $r(n)$ different solutions (some of which might be new solutions with v as the first vertex). At the end, the algorithm outputs the solution which contains the most vertices.

The multi-solution model is closely related to the advice complexity model. After processing the entire input, an algorithm in the multi-solution model has created at most $n \cdot r(n)$ different solutions (since at most $r(n)$ new solutions can be created in each round). Thus, one can convert a multi-solution algorithm to an algorithm with advice by letting the oracle provide $\log(n \cdot r(n))$ bits of advice indicating which solution to output. In addition, the oracle needs to provide $O(\log n)$ bits of advice in order to let the algorithm learn n (which was given to the multi-solution algorithm for free). On the other hand, an algorithm using $b(n)$ bits of advice can be converted to $2^{b(n)}$ deterministic algorithms. One can then run them in parallel to obtain a multi-solution algorithm with $r(n) = 2^{b(n)}$. These simple conversions allow one to translate both upper and lower bounds between the two models almost exactly (up to a lower-order additive term of $O(\log n)$).

It is shown in [19] that for any $c \geq 1$, there is a strictly c-competitive algorithm in the multi-solution model if $\lceil \log r(n) - 1 \rceil \geq n/c$. This gives a strictly c-competitive algorithm reading $\frac{n}{c} + O(\log n)$ bits of advice. On the other hand, it is shown that for any strictly c-competitive algorithm in the multi-solution model, it must hold that $c \geq n/(2\log(n \cdot r(n)))$. This implies that any strictly c-competitive algorithm with advice must read at least $\frac{n}{2c} - \log n$ bits of advice. Thus, the upper and lower bounds obtained in [19] are asymptotically tight.

Comparing our results to those of [19] and [7], we see that we improve on both the lower and upper bounds on the advice complexity of the problems under consideration by giving tight results. For the upper bound on ONLINE DISJOINT PATH ALLOCATION, the improvement is a factor of $(\log c)/2$. The results of [19] are already asymptotically tight. Our improvement consists of determining the exact coefficient of the higher-order term. Perhaps even more important, obtaining these tight lower and upper bounds on the advice complexity for ONLINE INDEPENDENT SET and ONLINE DISJOINT PATH ALLOCATION becomes very easy when using our string guessing problem ASG. We remark that the reductions we use to show the hardness of these problems reduces instances of ASG to instances of ONLINE INDEPENDENT SET (resp. ONLINE DISJOINT PATH ALLOCATION) that are identical to the hard instances used in [19] (resp. [7]). What enables us to improve the previous bounds, even though we use the same hard instances, is that we have a detailed analysis of the advice complexity of ASG at our disposal.

1.6 Related work

The advice complexity of ONLINE DISJOINT PATH ALLOCATION has also been studied as a function of the length of the path (as opposed to the number of requests), see [3, 7].

The advice complexity of ONLINE INDEPENDENT SET on bipartite graphs and on sparse graphs has been determined in [13]. It turns out that for these graph classes, even a small amount of advice can be very helpful. For instance, it is shown that a single bit of advice is enough to be 4-competitive on trees (recall that without advice, it is not possible to be better than $\Omega(n)$-competitive, even on trees).

It is clear that online maximum clique in the vertex arrival model is essentially equivalent to ONLINE INDEPENDENT SET. In [5], the advice complexity of a different version of online maximum clique is studied: The vertices of a graph are revealed as in the vertex-arrival model. Let V_{ALG} be the set of vertices selected by ALG and let C be a maximum clique in the subgraph induced by the vertices V_{ALG}. The profit of the solution V_{ALG} is $|C|^2 / |V_{\text{ALG}}|$. In particular, the algorithm is not required to output a clique, but is instead punished for including too many additional vertices in its output.

The ONLINE VERTEX COVER problem and some variations thereof are studied in [11].

The advice complexity of an online set cover problem [1] has been studied in [24]. However, the version of online set cover that we consider is different and so our results and those of [24] are incomparable.

2 Asymmetric String Guessing

In this section, we formally define the asymmetric string guessing problem and give simple algorithms for the problem. There are four variants of the problem, one for each combination of minimization/maximization and known/unknown history. Collectively, these four problems will be referred to as ASG.

We have deliberately tried to mimic the definition of the string guessing problem SG from [5]. However, for ASG, the number, n, of requests is not revealed to the online algorithm (as opposed to in [5]). This is only a minor technical detail since it changes the advice complexity by at most $O(\log n)$ bits.

2.1 The Minimization Version

We begin by defining the two minimization variants of ASG: One in which the output of the algorithm cannot depend on the correctness of previous answers (unknown history), and one in which the algorithm, after each guess, learns the correct answer (known history[1]). We collectively refer to the two minimization problems as MIN-ASG.

Definition 2 The *minimum asymmetric string guessing problem with unknown history*, MINASGU, has input $\langle ?_1, \ldots, ?_n, x \rangle$, where $x \in \{0,1\}^n$, for some $n \in \mathbb{N}$. For $1 \leq i \leq n$, round i proceeds as follows:

1. The algorithm receives request $?_i$ which contains no information.
2. The algorithm answers y_i, where $y_i \in \{0,1\}$.

The *output* $y = y_1 \ldots y_n$ computed by the algorithm is *feasible*, if $x \sqsubseteq y$. Otherwise, y is *infeasible*. The *cost* of a feasible output is $|y|_1$, and the cost of an infeasible output is ∞. The goal is to minimize the cost.

Thus, each request carries no information. While this may seem artificial, it does capture the hardness of some online problems (see for example Lemma 7).

Definition 3 The *minimum asymmetric string guessing problem with known history*, MINASGK, has input $\langle ?, x_1, \ldots, x_n \rangle$, where $x = x_1 \ldots x_n \in \{0,1\}^n$, for some $n \in \mathbb{N}$. For $1 \leq i \leq n$, round i proceeds as follows:

1. If $i > 1$, the algorithm learns the correct answer, x_{i-1}, to the request in the previous round.
2. The algorithm answers $y_i = f(x_1, \ldots, x_{i-1}) \in \{0,1\}$, where f is a function defined by the algorithm.

The *output* $y = y_1 \ldots y_n$ computed by the algorithm is *feasible*, if $x \sqsubseteq y$. Otherwise, y is *infeasible*. The *cost* of a feasible output is $|y|_1$, and the cost of an infeasible output is ∞. The goal is to minimize the cost.

[1] The concept of known history for online problems also appears in [19, 20] where it is denoted *transparency*.

The string x in either version of MINASG will be referred to as the *input string* or the *correct string*. Note that the number of requests in both versions of MINASG is $n+1$, since there is a final request that does not require any response from the algorithm. This final request ensures that the entire string x is eventually known. For simplicity, we will measure the advice complexity of MINASG as a function of n (this choice is not important as it changes the advice complexity by at most one bit).

Clearly, for any deterministic MINASG algorithm which sometimes answers 0, there exists an input string on which the algorithm gets a cost of ∞. However, if an algorithm always answers 1, the input string could consist solely of 0s. Thus, no deterministic algorithm can achieve any competitive ratio bounded by a function of n. One can easily show that the same holds for any randomized algorithm.

We now give a simple algorithm for MINASG which reads $O(n/c)$ bits of advice and achieves a strict competitive ratio of $\lceil c \rceil$.

Theorem 1 *For any $c \geq 1$, there is a strictly $\lceil c \rceil$-competitive algorithm for* MINASG *which reads $\lceil \frac{n}{c} \rceil + O(\log(n/c))$ bits of advice.*

Proof We will prove the result for MINASGU. Clearly, it then also holds for MINASGK.

Let $x = x_1 \dots x_n$ be the input string. The oracle encodes $p = \lceil n/c \rceil$ in a self-delimiting way, which requires $O(\log(n/c))$ bits of advice. For $0 \leq j < p$, define $C_j = \{x_i : i \equiv j \pmod{p}\}$. These p sets partition the input string, and the size of each C_j is at most $\lceil n/p \rceil$. The oracle writes one bit, b_j, for each set C_j. If C_j contains only 0s, b_j is set to 0. Otherwise, b_j is set to 1. Thus, in total, the oracle writes $\lceil n/c \rceil + O(\log(n/c))$ bits of advice to the advice tape.

The algorithm, ALG, learns p and the bits $b_0, \dots, b_{p-1}$ from the advice tape. In round i, ALG answers with the bit $b_{i \bmod p}$. We claim that this algorithm is strictly $\lceil c \rceil$-competitive. It is clear that the algorithm produces a feasible output. Furthermore, if ALG answers 1 in round i, it must be the case that at least one input bit in $C_{i \bmod p}$ is 1. Since the size of each C_j is at most $\lceil n/p \rceil \leq \lceil c \rceil$, this implies that ALG is strictly $\lceil c \rceil$-competitive. □

2.2 The Maximization Version

We also consider ASG in a maximization version. One can view this as a dual version of MINASG.

Definition 4 The *maximum asymmetric string guessing problem with unknown history*, MAXASGU, is identical to MINASGU, except that the score function is different: The score of a feasible output y is $|y|_0$, and the score of an infeasible output is $-\infty$. The goal is to maximize the score.

The maximum asymmetric string guessing problem with *known history* is defined similarly:

Definition 5 The *maximum asymmetric string guessing problem with known history*, MAXASGK, is identical to MINASGK, except that the score function is different: The

score of a feasible output y is $|y|_0$, and the score of an infeasible output is $-\infty$. The goal is to maximize the score.

We collectively refer to the two problems as MAXASG. Similarly, MINASGU and MAXASGU are collectively called ASGU, and MINASGK and MAXASGK are collectively called ASGK.

An algorithm for MAXASG without advice cannot attain any competitive ratio bounded by a function of n. If such an algorithm would ever answer 0 in some round, an adversary would let the correct answer be 1 and the algorithm's output would be infeasible. On the other hand, answering 1 in every round gives an output with a profit of zero.

Consider instances of MINASG and MAXASG with the same correct string x. It is clear that the optimal solution is the same for both instances. However, as is usual with dual versions of a problem, they differ with respect to approximation. For example, if half of the bits in x are 1s, then we get a 2-competitive solution y for the MINASG instance by answering 1 in each round. However, in MAXASG, the profit of the same solution y is zero. Despite this, there is a similar result to Theorem 1 for MAXASG.

Theorem 2 *For any $c \geq 1$, there is a strictly $\lceil c \rceil$-competitive algorithm for* MAXASG *which reads $\lceil n/c \rceil + O(\log n)$ bits of advice.*

Proof We will prove the result for MAXASGU. Clearly, it then also holds for MAXASGK.

The oracle partitions the input string $x = x_1 \dots x_n$ into $\lceil c \rceil$ disjoint blocks, each containing (at most) $\lceil \frac{n}{c} \rceil$ consecutive bits. Note that there must exist a block where the number of 0s is at least $|x|_0 / \lceil c \rceil$. The oracle uses $O(\log n)$ bits to encode the index i in which this block starts and the index i' in which it ends. Furthermore, the oracle writes the string $x_i \dots x_{i'}$ onto the advice tape, which requires at most $\lceil \frac{n}{c} \rceil$ bits, since this is the largest possible size of a block. The algorithm learns the string $x_i \dots x_{i'}$ and answers accordingly in rounds i to i'. In all other rounds, the algorithm answers 1. Since the profit of this output is at least $|x|_0 / \lceil c \rceil$, it follows that ALG is strictly $\lceil c \rceil$-competitive. □

In the following section, we determine the amount of advice an algorithm needs to achieve some competitive ratio $c > 1$. It turns out that the algorithms from Theorems 1 and 2 use the asymptotically smallest possible number of advice bits, but the coefficient in front of the term n/c can be improved.

3 Advice Complexity of ASG

In this section we give upper and lower bounds on the number of advice bits necessary to obtain c-competitive ASG algorithms, for some $c > 1$. The bounds are tight up to $O(\log n)$ bits. For ASGU, the gap between the upper and lower bounds stems only from the fact that the advice used for the upper bound includes the number, n, of requests and the number, t, of 1-bits in the input. Since the lower bound is shown to hold even if the algorithm knows n and t, this slight gap is to be expected.

The following two observations will be used extensively in the analysis.

Remark 1 Suppose that a MINASG algorithm, ALG, is c-competitive. By definition, there exists a constant, α, such that $\text{ALG}(\sigma) \leq c \cdot \text{OPT}(\sigma) + \alpha$. Then, one can construct a new algorithm, ALG$'$, which is *strictly* c-competitive and uses $O(\log n)$ additional advice bits as follows:

Use $O(\log n)$ bits of advice to encode the length n of the input and use $\alpha \cdot \lceil \log n \rceil = O(\log n)$ bits of advice to encode the index of (at most) α rounds in which ALG guesses 1 but where the correct answer is 0. Clearly, ALG$'$ can use this additional advice to achieve a strict competitive ratio of c.

This also means that a lower bound of b on the number of advice bits required to be *strictly* c-competitive implies a lower bound of $b - O(\log n)$ advice bits for being c-competitive (where the constant hidden in $O(\log n)$ depends on the additive constant α of the c-competitive algorithm).

The same technique can be used for MAXASG.

Remark 2 For a minimization problem, an algorithm, ALG, using b bits of advice can be converted into 2^b algorithms, $\text{ALG}_1, \ldots, \text{ALG}_{2^b}$, without advice, one for each possible advice string, such that $\text{ALG}(\sigma) = \min_i \text{ALG}_i(\sigma)$ for any input sequence σ. The same holds for maximization problems, except that in this case, $\text{ALG}(\sigma) = \max_i \text{ALG}_i(\sigma)$.

For ASG with unknown history, the output of a deterministic algorithm can depend only on the advice, since no information is revealed to the algorithm through the input. Thus, for MINASGU and MAXASGU, a deterministic algorithm using b advice bits can produce only 2^b different outputs, one for each possible advice string.

3.1 Using Covering Designs

In order to determine the advice complexity of ASG, we will use some basic results from the theory of combinatorial designs. We start with the definition of a covering design.

For any $k \in \mathbb{N}$, a *k-set* is a set of cardinality k. Let $v \geq k \geq t$ be positive integers. A *(v,k,t)-covering design* is a family of k-subsets (called *blocks*) of a v-set, S, such that any t-subset of S is contained in at least one block. The *size* of a covering design, D, is the number of blocks in D. The *covering number*, $C(v,k,t)$, is the smallest possible size of a (v,k,t)-covering design. Many papers have been devoted to the study of these numbers. See [12] for a survey. The connection to ASG is that for inputs to MINASG where the number of 1s is t, an $(n, \lfloor ct \rfloor, t)$-covering design can be used to obtain a strictly c-competitive algorithm.

It is clear that a (v,k,t)-covering design always exists. Since a single block has exactly $\binom{k}{t}$ t-subsets, and since the total number of t-subsets of a set of size v is $\binom{v}{t}$, it follows that $\binom{v}{t} / \binom{k}{t} \leq C(v,k,t)$. We will make use of the following upper bound on the size of a covering design:

Lemma 1 (Erdős, Spencer [16]) *For all natural numbers $v \geq k \geq t$,*

$$\frac{\binom{v}{t}}{\binom{k}{t}} \leq C(v,k,t) \leq \frac{\binom{v}{t}}{\binom{k}{t}} \left(1 + \ln \binom{k}{t}\right)$$

We use Lemma 1 to express both the upper and lower bound in terms of (a quotient of) binomial coefficients. This introduces an additional difference of $\log n$ between the stated lower and upper bounds.

Lemma 17 in Appendix A shows how the bounds we obtain can be approximated by a closed formula, avoiding binomial coefficients. This approximation costs an additional (additive) difference of $O(\log n)$ between the lower and upper bounds. The approximation is in terms of the following function:

$$B(n,c) = \log\left(1 + \frac{(c-1)^{c-1}}{c^c}\right) n$$

For $c > 1$, we show that $B(n,c) \pm O(\log n)$ bits of advice are necessary and sufficient to achieve a (strict) competitive ratio of c, for any version of ASG. See Figure 1 for a graphical view. It can be shown (Lemma 15) that

$$\frac{1}{e\ln(2)}\frac{n}{c} \leq B(n,c) \leq \frac{n}{c}.$$

In particular, if $c = o(n/\log n)$, we see that $O(\log n)$ becomes a lower-order additive term. Thus, for this range of c, we determine exactly the higher-order term in the advice complexity of ASG. Since this is the main focus of our paper, we will often refer to $O(\log n)$ as a lower-order additive term. The case where $c = \Omega(n/\log n)$ is treated separately in Section 3.4.

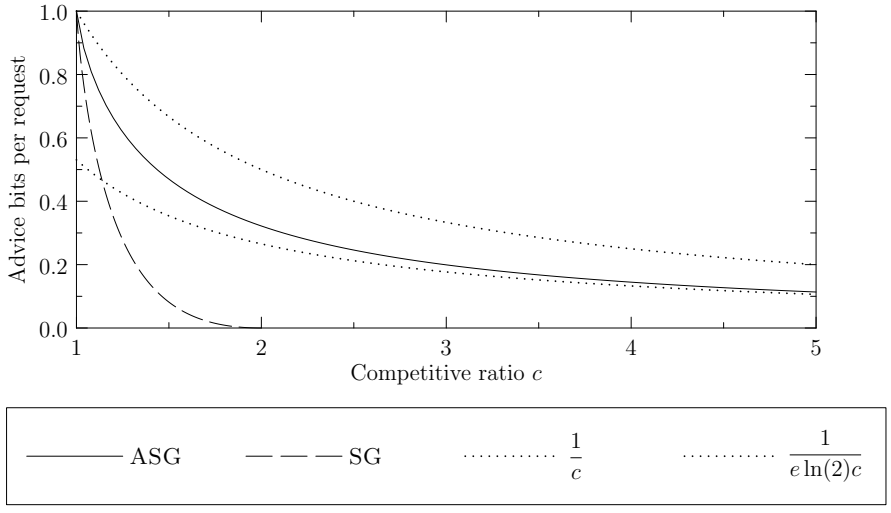

Fig. 1 The solid line shows the number of advice bits per request which are necessary and sufficient for obtaining a (strict) competitive ratio of c for ASG (ignoring lower-order terms). The dashed line shows the same number for the original binary string guessing problem SG [5]. The dotted lines are the functions $1/c$ and $1/(e\ln(2)c)$.

3.2 Advice Complexity of MINASG

We first consider MINASG with unknown history. Clearly, an upper bound for MINASGU is also valid for MINASGK. We will show that the covering number $C(v,k,t)$ is very closely related to the advice complexity of MINASGU.

Theorem 3 *For any $c > 1$, there exists a strictly c-competitive algorithm for* MINASG *reading b bits of advice, where*

$$b \leq B(n,c) + O(\log n).$$

Proof We will define an algorithm ALG and an oracle $\mathtt{O}$ for MINASGU such that ALG is strictly c-competitive and reads at most b bits of advice. Clearly, the same algorithm can be used for MINASGK.

Let $x = x_1 \dots x_n$ be an input string to MINASGU and set $t = |x|_1$. The oracle $\mathtt{O}$ writes the value of n to the advice tape using a self-delimiting encoding. Furthermore, the oracle writes the value of t to the advice tape using $\lceil \log n \rceil$ bits (this is possible since $t \leq n$). Thus, this part of the advice uses at most $3\lceil \log n \rceil + 1$ bits in total.

If $\lfloor ct \rfloor \geq n$, then ALG will answer 1 in each round. If $t = 0$, ALG will answer 0 in each round.

If $0 < \lfloor ct \rfloor < n$, then ALG computes an optimal $(n, \lfloor ct \rfloor, t)$-covering design as follows: ALG tries (in lexicographic order, say) all possible sets of $\lfloor ct \rfloor$-blocks, starting with sets consisting of one block, then two blocks, and so on. For each such set, ALG can check if it is indeed an $(n, \lfloor ct \rfloor, t)$-covering design. As soon as a valid covering design, D, is found, the algorithm can stop, since D will be a smallest possible $(n, \lfloor ct \rfloor, t)$-covering design.

Now, $\mathtt{O}$ picks a $\lfloor ct \rfloor$-block, S_y, from D, such that the characteristic vector y of S_y satisfies that $x \sqsubseteq y$. Note that, since ALG is deterministic, the oracle knows which covering design ALG computes and the ordering of the blocks in that design. The oracle then writes the index of S_y on the advice tape. This requires at most $\lceil \log C(n, \lfloor ct \rfloor, t) \rceil$ bits of advice.

ALG reads the index of the $\lfloor ct \rfloor$-block S_y from the advice tape and answers 1 in round i if and only if the element i belongs to S_y. Clearly, this will result in ALG answering 1 exactly $\lfloor ct \rfloor$ times and producing a feasible output. It follows that ALG is strictly c-competitive. Furthermore, the number of bits read by ALG is

$$b \leq \left\lceil \log \left(\max_{t \colon \lfloor ct \rfloor < n} C(n, \lfloor ct \rfloor, t) \right) \right\rceil + 3\lceil \log n \rceil + 1 \,.$$

The theorem now follows from Lemma 17, Inequality (8). □

We now give an almost matching lower bound.

Theorem 4 *For any $c > 1$, a c-competitive algorithm* ALG *for* MINASGU *must read b bits of advice, where*

$$b \geq B(n,c) - O(\log n) \,.$$

Proof By Remark 1, it suffices to prove the lower bound for strictly c-competitive algorithms. Suppose that ALG is strictly c-competitive. Let b be the number of advice bits read by ALG on inputs of length n. For $0 \leq t \leq n$, let $I_{n,t}$ be the set of input strings of length n with Hamming weight t, and let $Y_{n,t}$ be the corresponding set of output strings produced by ALG. We will argue that, for each t, $0 \leq \lfloor ct \rfloor \leq n$, $Y_{n,t}$ can be converted to an $(n, \lfloor ct \rfloor, t)$-covering design of size at most 2^b.

By Remark 2, ALG can produce at most 2^b different output strings, one for each possible advice string. Now, for each input string, $x \in I_{n,t}$, there must exist some advice which makes ALG output a string y, where $|y|_1 \leq \lfloor ct \rfloor$ and $x \sqsubseteq y$. If not, then ALG is not strictly c-competitive. For each possible output $y \in \{0,1\}^n$ computed by ALG, we convert it to the set $S_y \subseteq [n]$ which has y as its characteristic vector. If $|y|_1 < \lfloor ct \rfloor$, we add some arbitrary elements to S_y so that S_y contains exactly $\lfloor ct \rfloor$ elements. Since ALG is strictly c-competitive, this conversion gives the blocks of an $(n, \lfloor ct \rfloor, t)$-covering design. The size of this covering design is at most 2^b, since ALG can produce at most 2^b different outputs. It follows that $C(n, \lfloor ct \rfloor, t) \leq 2^b$, for all t, $0 \leq \lfloor ct \rfloor \leq n$. Thus,

$$b \geq \log \left(\max_{t\colon \lfloor ct \rfloor < n} C(n, \lfloor ct \rfloor, t) \right).$$

The theorem now follows from Lemma 17, Inequality (6). □

Note that the proof of Theorem 4 relies heavily on the unknown history in order to bound the total number of possible outputs. However, Theorem 5 below states that the lower bound of $B(n,c) - O(\log n)$ also holds for MINASGK. In order to prove this, we show how an adversary can ensure that revealing the correct answers for previous requests does not give the algorithm too much extra information. The way to ensure this depends on the specific strategy used by the algorithm and oracle at hand, and so the proof is more complicated than that of Theorem 4.

Theorem 5 *For any $c > 1$, a c-competitive algorithm for* MINASGK *must read b bits of advice, where*

$$b \geq B(n,c) - O(\log n).$$

Proof By Remark 1, it suffices to prove the lower bound for strictly c-competitive algorithms. Consider the set, $I_{n,t}$, of input strings of length n and Hamming weight t, for some t such that $\lfloor ct \rfloor \leq n$. Restricting the input set to strings with one particular Hamming weight can only weaken the adversary.

Let ALG be a strictly c-competitive algorithm for MINASGK which reads at most b bits of advice for any input of length n. For an advice string φ, denote by $I_\varphi \subseteq I_{n,t}$ the set of input strings for which ALG reads the advice φ. Since we are considering MINASGK, in any round, ALG may use both the advice string and the information about the correct answer for previous rounds when deciding on an answer for the current round.

We will prove the lower bound by considering the computation of ALG, when reading the advice φ, as a game between ALG and an adversary. This game proceeds according to the rules specified in Definition 3. In particular, at the beginning of round

i, the adversary reveals the correct answer x_{i-1} for round $i-1$ to ALG. Thus, at the beginning of round i, the algorithm knows the first $i-1$ bits, $x_1,\ldots,x_{i-1}$, of the input string. We say that a string $s \in I_\varphi$ is *alive* in round i if $s_j = x_j$ for all $j < i$, and we denote by $I_\varphi^i \subseteq I_\varphi$ the set of strings which are alive in round i. The adversary must reveal the correct answers in a way that is consistent with φ. That is, in each round, there must exist at least one string in I_φ which is alive.

We first make two simple observations:

- Suppose that, in some round i, there exists a string $s \in I_\varphi^i$ such that $s_i = 1$. Then, ALG must answer 1, or else the adversary can choose s as the input string and thereby force ALG to incur a cost of ∞. Thus, we will assume that ALG always answers 1 in such rounds.
- On the other hand, if, in round i, all $s \in I_\varphi^i$ have $s_i = 0$, then ALG is free to answer 0. We will assume that ALG always answers 0 in such rounds.

Assume that, at some point during the computation, I_φ contains exactly m strings and exactly h 1s are still to be revealed. We let $L_1(m,h)$ be the largest number such that for every set of m different strings of equal length, each with Hamming weight h, the adversary can force ALG to incur a cost of at least $L_1(m,h)$ when starting for this situation. In other words, $L_1(m,h)$ is the minimum number of rounds in which the adversary can force ALG to answer 1.

Claim: For any $m,h \geq 1$,

$$L_1(m,h) \geq \min\left\{d\colon m \leq \binom{d}{h}\right\}. \tag{2}$$

Before proving the claim, we will show how it implies the theorem. For any t, $0 \leq \lfloor ct \rfloor < n$, there are $\binom{n}{t}$ possible input strings of length n and Hamming weight t. By the pigeonhole principle, there must exist an advice string φ' such that $\left|I_{\varphi'}\right| \geq \binom{n}{t}/2^b$. Now, if $m = \left|I_{\varphi'}\right| > \binom{\lfloor ct \rfloor}{t}$, then by (2), $L_1(m,t) \geq \min\{d\colon \binom{\lfloor ct \rfloor}{t} < \binom{d}{t}\} = \lfloor ct \rfloor + 1$. This contradicts the fact that ALG is strictly c-competitive. Thus, it must hold that $\left|I_{\varphi'}\right| \leq \binom{\lfloor ct \rfloor}{t}$. Combining the two inequalities involving $\left|I_{\varphi'}\right|$, we get

$$\binom{\lfloor ct \rfloor}{t} \geq \left|I_{\varphi'}\right| \geq \frac{\binom{n}{t}}{2^b} \;\Rightarrow\; 2^b \geq \frac{\binom{n}{t}}{\binom{\lfloor ct \rfloor}{t}}$$

Since this holds for all values of t, we obtain the lower bound

$$b \geq \log\left(\max_{t\colon \lfloor ct \rfloor < n} \frac{\binom{n}{t}}{\binom{\lfloor ct \rfloor}{t}}\right).$$

The theorem then follows from Lemma 17 and Inequalities (7) and (6).

Proof of claim: Fix $1 \leq i \leq n$ and assume that, at the beginning of round i, there are m strings alive, all of which still have exactly h 1's to be revealed. The rest of the proof is by induction on m and h.

For the *base case*, suppose first that $h = 1$. Then, for each of the m strings, $s^1,\ldots,s^m \in I_\varphi^i$, there is exactly one index, $i_1,\ldots,i_m$, such that $s^1_{i_1} = \cdots = s^m_{i_m} = 1$.

Since all strings in I_φ^i must be different, it follows that $i_j \neq i_k$ for $j \neq k$. Without loss of generality, assume that $i_1 < i_2 < \cdots < i_m$. In rounds $i_1, \ldots, i_{m-1}$, the adversary chooses the correct answer to be 0, while ALG is forced to answer 1 in each of these rounds. Finally, in round i_m, the adversary reveals the correct answer to be 1 (and hence the input string must be s^m). In total, ALG incurs a cost of m, which shows that $L_1(m,1) = m$ for all $m \geq 1$.

Assume now that $m = 1$. It is clear that $L_1(m,h) \geq h$ for all values of h. In particular, $L_1(1,h) = h$. This finishes the base case.

For the *inductive step*, fix integers $m,h \geq 2$. Assume that the formula is true for all (i,j) such that $j \leq h-1$ or such that $j = h$ and $i \leq m-1$. We will show that the formula is also true for (m,h).

Consider the strings $s^1, \ldots, s^m \in I_\varphi^i$ alive at the beginning of round i. We partition I_φ^i into two sets, $S_0 = \{s^j \colon s_i^j = 0\}$ and $S_1 = \{s^j \colon s_i^j = 1\}$, and let $m_0 = |S_0|$ and $m_1 = |S_1|$. Recall that if all sequences $s \in I_\varphi^i$ have $s_i = 0$, we assume that ALG answers 0, leaving m and h unchanged. Thus, we may safely ignore such rounds and assume that $m_0 < m$. We let

$$d = \min\left\{d' \colon m \leq \binom{d'}{h}\right\},$$
$$d_0 = \min\left\{d' \colon m_0 \leq \binom{d'}{h}\right\}, \text{ and}$$
$$d_1 = \min\left\{d' \colon m_1 \leq \binom{d'}{h-1}\right\}.$$

If $d_1 + 1 \geq d$, then the adversary chooses 1 as the correct answer in round i. By the induction hypothesis, $L_1(m_1, h-1) \geq d_1$. Together with the fact that ALG is forced to answer 1 in round i, this shows that the adversary can force ALG to incur a cost of at least $L_1(m_1, h-1) + 1 \geq d_1 + 1 \geq d$.

On the other hand, if $d_1 + 1 < d$, the adversary chooses 0 as the correct answer in round i. Note that this implies that each string alive in round $i+1$ still has exactly h 1's to be revealed. We must have $d_1 \leq d-2$ since d_1 and d are both integers. Moreover, by definition of d, it holds that $m > \binom{d-1}{h}$. Thus, we get the following lower bound on m_0:

$$\begin{aligned}
m_0 &= m - m_1 \\
&> \binom{d-1}{h} - \binom{d_1}{h-1} \\
&\geq \binom{d-1}{h} - \binom{d-2}{h-1}, \text{ since } \binom{a}{b} \text{ is increasing in } a \\
&= \binom{d-2}{h}, \text{ by Pascal's Identity.}
\end{aligned}$$

This lower bound on m_0 shows that $d_0 > d-2$, and hence $d_0 \geq d-1$. Combining this with the induction hypothesis gives $L_1(m_0,h) \geq d_0 \geq d-1$. Since $m_1 \geq 1$, ALG is forced to answer 1 in round i, so the adversary can make ALG incur a cost of at least $L_1(m_0,h) + 1 \geq d$. □

3.3 Advice Complexity of MAXASG

In this section, we will show that the advice complexity of MAXASG is the same as that of MINASG, up to a lower-order additive term of $O(\log n)$. We use the same techniques as in Section 3.2.

As noted before, the difficulty of computing a c-competitive solution for a specific input string is not the same for MINASG and MAXASG. The key point is that computing a c-competitive solution for MAXASG, on input strings with u 0's, is roughly as difficult as computing a c-competitive solution for MINASG, on input strings with $\lceil u/c \rceil$ 1's.

We show that the proofs of Theorems 3–5 can easily be modified to give upper and lower bounds on the advice complexity of MAXASG. These bounds within the proofs look slightly different from the ones obtained for MINASG, but we show in Lemmas 19 and 20 that they differ from $B(n,c)$ by at most an additive term of $O(\log n)$.

Theorem 6 *For any $c > 1$, there exists a strictly c-competitive online algorithm for* MAXASG *reading b bits of advice, where*

$$b \leq B(n,c) + O(\log n).$$

Proof We will define an algorithm ALG and an oracle O for MAXASGU such that ALG is strictly c-competitive and reads at most b bits of advice. Clearly, the same algorithm can be used for MAXASGK.

As in the proof of Theorem 3, we note that, for any integers n, u where $0 < u < n$, the algorithm ALG can compute an optimal $(n, n - \lceil u/c \rceil, n-u)$-covering design deterministically.

Let $x = x_1 \dots x_n$ be an input string to MAXASGU and set $u = |x|_0$. The oracle O writes the values of n and u to the advice tape using at most $3\lceil \log n \rceil + 1$ bits in total.

If $0 < u < n$, then O picks an $(n - \lceil u/c \rceil)$-block, S_y, from the optimal $(n, n - \lceil u/c \rceil, n-u)$-covering design, as computed by ALG, such that the characteristic vector y of S_y satisfies that $x \sqsubseteq y$. The oracle writes the index of S_y on the advice tape. This requires at most $\lceil \log C(n, n - \lceil u/c \rceil, n-u) \rceil$ bits of advice.

The algorithm, ALG, first reads the values of n and u from the advice tape. If $u = 0$, then ALG will answer 1 in each round, and if $u = n$, then ALG will answer 0 in each round. If $0 < u < n$, then ALG will read the index of the $(n - \lceil u/c \rceil)$-block S_y from the advice tape. ALG will answer 1 in round i if and only if the element i belongs to the given block. Clearly, this will result in ALG answering 0 exactly $n - (n - \lceil u/c \rceil) = \lceil u/c \rceil$ times and producing a feasible output. It follows that ALG will be strictly c-competitive. Furthermore, the number of bits read by ALG is

$$b \leq \left\lceil \log \left(\max_{u:\ 0<u<n} C\left(n, n - \left\lceil \frac{u}{c} \right\rceil, n-u\right) \right) \right\rceil + 3\lceil \log n \rceil + 1 \,.$$

The theorem now follows from Lemma 20. □

Theorem 7 *For any $c > 1$, a c-competitive algorithm* ALG *for* MAXASGU *must read b bits of advice, where*

$$b \geq B(n,c) - O(\log n) \,.$$

Proof By Remark 1, it suffices to prove the lower bound for strictly c-competitive algorithms. Suppose that ALG is strictly c-competitive. Let b be the number of advice bits read by ALG on inputs of length n. For $0 \leq u \leq n$, let $I_{n,u}$ be the set of input strings x of length n with $|x|_0 = u$, and let $Y_{n,u}$ be the corresponding set of output strings produced by ALG. We will argue that, for each u, $Y_{n,u}$ can be converted to an $(n, n - \lceil u/c \rceil, n-u)$-covering design of size at most 2^b.

By Remark 2, ALG can produce at most 2^b different output strings, one for each possible advice string. Now, for each input string, $x = x_1 \ldots x_n$ with $|x|_0 = u$ (and, hence, $|x|_1 = n-u$), there must exist some advice which makes ALG output a string $y = y_1 \ldots y_n$ where $|y|_0 \geq \lceil u/c \rceil$ (and, hence, $|y|_1 \leq n - \lceil u/c \rceil$) and $x \sqsubseteq y$. If not, then ALG is not strictly c-competitive. For each possible output $y \in \{0,1\}^n$ computed by ALG, we convert it to the set $S_y \subseteq [n]$ which has y as its characteristic vector. If $|y|_1 < n - \lceil u/c \rceil$, we add some arbitrary elements to S_y so that S_y contains exactly $n - \lceil u/c \rceil$ elements. Since ALG is strictly c-competitive, this conversion gives the blocks of an $(n, n - \lceil u/c \rceil, n-u)$-covering design. The size of this covering design is at most 2^b, since ALG can produce at most 2^b different outputs. It follows that $C(n, n - \lceil u/c \rceil, n-u) \leq 2^b$, for all u. Thus,

$$b \geq \log \left(\max_{u:\, 0<u<n} C\left(n, n - \left\lceil \frac{u}{c} \right\rceil, n-u\right) \right).$$

The theorem now follows from Lemma 20. □

As was the case for MINASG, the lower bound for MAXASGU also holds for MAXASGK.

Theorem 8 *For any $c > 1$, a c-competitive algorithm* ALG *for* MAXASGK *must read at least b bits of advice, where*

$$b \geq B(n,c) - O(\log n).$$

Proof By Remark 1, it suffices to prove the lower bound for strictly c-competitive algorithms.

Consider input strings, x, of length n and such that $|x|_0 = u$. Let $t = |x|_1 = n-u$. We reuse the notation from the proof of Theorem 5 and let $I_\varphi \subseteq I_{n,t}$ denote the set of strings for which ALG reads the advice string φ.

Suppose there exists some advice string φ' such that $m = |I_{\varphi'}| > \binom{n - \lceil \frac{u}{c} \rceil}{t}$. Since Inequality (2) from the proof of Theorem 5 holds for MAXASG too, we get that $L_1(m,t) \geq n - \lceil \frac{u}{c} \rceil + 1$. But this means that there exists an input $x \in I_{\varphi'}$, with $|x|_1 = t$, such that ALG must answer 1 at least $n - \lceil \frac{u}{c} \rceil + 1$ times. In other words, for the output y, computed by ALG on input x, it holds that $|y|_0 \leq n - (n - \lceil \frac{u}{c} \rceil + 1) \leq \lceil \frac{u}{c} \rceil - 1$. Since $|x|_0 = u$, this contradicts the fact that ALG is strictly c-competitive.

Since there are $\binom{n}{u}$ possible input strings x such that $|x|_0 = u$, and since the above was shown to hold for all choices of u, we get the lower bound

$$b \geq \log \left(\max_{u:\, 0<u<n} \frac{\binom{n}{u}}{\binom{n - \lceil \frac{u}{c} \rceil}{n-u}} \right).$$

The theorem now follows from Lemma 20. □

3.4 Advice Complexity of ASG when $c = \Omega(n/\log n)$

Throughout the paper, we mostly ignore additive terms of $O(\log n)$ in the advice complexity. However, in this section, we will consider the advice complexity of ASG when the number of advice bits read is at most logarithmic. Surprisingly, it turns out that the advice complexity of MINASG and MAXASG is different in this case.

Recall that, by Theorem 6 (or Theorem 2), using $O(\log n)$ bits of advice, an algorithm for MAXASG can achieve a competitive ratio of $\frac{n}{\log n}$. The following theorem shows that there is a "phase-transition" in the advice complexity, in the sense that using less than $\log n$ bits of advice is no better than using no advice at all. We remark that Theorem 9 and its proof are essentially equivalent to a previous result of Halldórsson et al. [19] on ONLINE INDEPENDENT SET in the multi-solution model.

Theorem 9 (cf. [19]) *Let* ALG *be an algorithm for* MAXASG *reading* $b < \lfloor \log n \rfloor$ *bits of advice. Then, the competitive ratio of* ALG *is not bounded by a function of n. This is true even if* ALG *knows n in advance.*

Proof We will prove the result for MAXASGK. Clearly, it then also holds for MAXASGU.

By Remark 2, we can convert ALG to $m = 2^b$ online algorithms without advice. Denote the algorithms by $\text{ALG}_1, \ldots, \text{ALG}_m$. Since $b < \lfloor \log n \rfloor$, it follows that $m \leq n/2$. We claim that the adversary can construct an input string $x = x_1 \ldots x_n$ for MAXASGK such that the following holds: For each $1 \leq j \leq m$, the output of ALG_j is either infeasible or contains only 1s. Furthermore, x can be constructed such that $|x|_0 \geq \frac{n}{2}$.

We now show how the adversary may achieve this. For $1 \leq i \leq n$, the adversary decides the value of x_i as follows: If there is some algorithm, ALG_j, which answers 0 in round i and ALG_j answers 1 in all rounds before round i, the adversary lets $x_i = 1$. In all other cases, the adversary lets $x_i = 0$. It follows that if an algorithm ALG_j ever answers 0, its output will be infeasible. Furthermore, the number of 1's in the input string constructed by the adversary is at most $n/2$, since $m \leq n/2$. Thus, the profit of OPT on this input is at least $n/2$, while the profit of ALG is at most 0. □

For MINASG, the algorithm from Theorem 1 achieves a competitive ratio of $\lceil c \rceil$ and uses $O(n/c)$ bits of advice, for any $c > 1$. In particular, it is possible to achieve a competitive ratio of e.g. $O(n/(\log\log n))$ using $O(\log\log n)$ bits of advice, which we have just shown is not possible for MAXASG. The following theorem shows that no strictly c-competitive algorithm for MINASG can use less than $\Omega(n/c)$ bits of advice, even if $n/c = o(\log n)$.

Theorem 10 *For any* $c > 1$*, on inputs of length n, a strictly* $\lceil c \rceil$*-competitive algorithm* ALG *for* MINASG *must read at least* $b = \Omega(n/c)$ *bits of advice.*

Proof We will prove the result for MINASGK. Clearly, it then also holds for MINASGU.

Suppose that ALG is strictly $\lceil c \rceil$-competitive. Since $\lfloor \lceil c \rceil t \rfloor = \lceil c \rceil t$, it follows from the proof of Theorem 5 that ALG must read at least b bits of advice, where

$$b \geq \log \left(\max_{t:\, \lceil c \rceil t < n} \frac{\binom{n}{t}}{\binom{\lceil c \rceil t}{t}} \right).$$

By Lemma 18, this implies that $b = \Omega(n/c)$. □

4 The Complexity Class AOC

In this section, we define a class, AOC, and show that for each problem, P, in AOC, the advice complexity of P is at most that of ASG.

Definition 6 A problem, P, is in AOC *(Asymmetric Online Covering)* if it can be defined as follows: The input to an instance of P consists of a sequence of n requests, $\sigma = \langle r_1, \ldots, r_n \rangle$, and possibly one final dummy request. An algorithm for P computes a binary output string, $y = y_1 \ldots y_n \in \{0,1\}^n$, where $y_i = f(r_1, \ldots, r_i)$ for some function f.

For minimization (maximization) problems, the score function, s, maps a pair, (σ, y), of input and output to a cost (profit) in $\mathbb{N} \cup \{\infty\}$ ($\mathbb{N} \cup \{-\infty\}$). For an input, σ, and an output, y, y is *feasible* if $s(\sigma, y) \in \mathbb{N}$. Otherwise, y is *infeasible*. There must exist at least one feasible output. Let $S_{\min}(\sigma)$ ($S_{\max}(\sigma)$) be the set of those outputs that minimize (maximize) s for a given input σ.

If P is a minimization problem, then for every input, σ, the following must hold:

1. For a feasible output, y, $s(\sigma, y) = |y|_1$.
2. An output, y, is feasible if there exists a $y' \in S_{\min}(\sigma)$ such that $y' \sqsubseteq y$.
   If there is no such y', the output may or may not be feasible.

If P is a maximization problem, then for every input, σ, the following must hold:

1. For a feasible output, y, $s(\sigma, y) = |y|_0$.
2. An output, y, is feasible if there exists a $y' \in S_{\max}(\sigma)$ such that $y' \sqsubseteq y$.
   If there is no such y', the output may or may not be feasible.

The dummy request is a request that does not require an answer and is not counted when we count the number of requests. Most of the problems that we consider will not have such a dummy request, but it is necessary to make sure that ASG belongs to AOC.

The input, σ, to a problem P in AOC can contain any kind of information. However, for each request, an algorithm for P only needs to make a binary decision. If the problem is a minimization problem, it is useful to think of answering 1 as accepting the request and answering 0 as rejecting the request (e.g. vertices in a vertex cover). The output is guaranteed to be feasible if the accepted requests are a superset of the requests accepted in an optimal solution (they "cover" the optimal solution).

If the problem is a maximization problem, it is useful to think of answering 0 as accepting the request and answering 1 as rejecting the request (e.g. vertices in an

independent set). The output is guaranteed to be feasible if the accepted requests are a subset of the requests accepted in a optimal solution.

Note that outputs for problems in AOC may have a score of $\pm\infty$. This is used to model that the output is infeasible (e.g. not a vertex cover/independent set).

We now show that our ASGU algorithm based on covering designs works for every problem in AOC. This gives an upper bound on the advice complexity for all problems in AOC.

Theorem 11 *Let* P *be a problem in* AOC. *There exists a strictly c-competitive online algorithm for* P *reading b bits of advice, where*

$$b \leq B(n,c) + O(\log n).$$

Proof We first assume that P is a minimization problem. Let ALG be a strictly c-competitive MINASGU algorithm reading at most b bits of advice provided by an oracle O. By Theorem 3, such an algorithm exists. We will define a P algorithm, ALG$'$, together with an oracle O', that is strictly c-competitive and reads at most b bits of advice.

For a given input, σ, to P, the oracle O' starts by computing an x such that $x \in S_{\min}(\sigma)$. This is always possible since by the definition of AOC, such an x always exists, and O' has unlimited computational power. Let φ be the advice that O would write to the advice tape if x were the input string in an instance of MINASGU. O' writes φ to the advice tape. From here, ALG$'$ behaves as ALG would do when reading φ (in particular, ALG$'$ ignores any possible information contained in σ) and computes the output y. Since ALG is strictly c-competitive for MINASGU, we know that $x \sqsubseteq y$ and that $|y|_1 \leq c\,|x|_1$. Since P is in AOC, this implies that y is feasible (with respect to the input σ) and that $s(\sigma,y) \leq c\,|x|_1 = c \cdot \textsc{Opt}(\sigma)$.

Similarly, one can reduce a maximization problem to MAXASGU and apply Theorem 6. □

Showing that a problem, P, belongs to AOC immediately gives an upper bound on the advice complexity of P. For all variants of ASG, we know that this upper bound is tight up to an additive $O(\log n)$ term. This leads us to the following definition of completeness.

Definition 7 A problem, P, is AOC-*complete* if

- P belongs to AOC and
- for all $c > 1$, any c-competitive algorithm for P must read at least b bits of advice, where

$$b \geq B(n,c) - O(\log n).$$

Thus, the advice complexity of an AOC-complete problem must be identical to the upper bound from Theorem 11, up to a lower-order additive term of $O(\log n)$. By Definitions 2–5 combined with Theorems 4–5 and 7–8, all of MINASGU, MINASGK, MAXASGU and MAXASGK are AOC-complete.

When we show that some problem, P, is AOC-complete, we usually do this by giving a reduction from a known AOC-complete problem to P, preserving the competitive ratio and increasing the number of advice bits by at most $O(\log n)$. ASGK is especially well-suited as a starting point for such reductions.

We allow for an additional $O(\log n)$ bits of advice in Definition 7 in order to be able to use the reduction between the strict and non-strict competitive ratios as explained in Remark 1 and in order to encode some natural parameters of the problem, such as the input length or the score of an optimal solution. For most values of c, it seems reasonable to allow these additional advice bits. However, it does mean that for $c = \Omega(n/\log n)$, the requirement in the definition of AOC-complete is vacuously true. We refer to Section 3.4 for a discussion of the advice complexity for this range of competitive ratio.

4.1 AOC-complete Minimization Problems

In this section, we show that several online problems are AOC-complete, starting with ONLINE VERTEX COVER. See the introduction for the definition of the problems

4.1.1 Online Vertex Cover.

Lemma 2 ONLINE VERTEX COVER *is in* AOC.

Proof We need to verify the conditions in Definition 6.

Recall that an input $\sigma = \langle r_1, \ldots, r_n \rangle$ for ONLINE VERTEX COVER is a sequence of requests, where each request is a vertex along with the edges connecting it to previously requested vertices. There is no dummy request at the end. For each request, r_i, an algorithm makes a binary choice, y_i: It either includes the vertex into its solution ($y_i = 1$) or not ($y_i = 0$).

The cost of an infeasible solution is ∞. A solution $y = y_1 \ldots y_n$ for ONLINE VERTEX COVER is feasible if the vertices included in the solution form a vertex cover in the input graph. Clearly, there is always at least one feasible solution, since taking all the vertices will give a vertex cover.

Thus, ONLINE VERTEX COVER has the right form. Finally, we verify that conditions 1 and 2 are also satisfied: Condition 1 is satisfied since the cost of a feasible solution is the number of vertices in the solution and condition 2 is satisfied since a superset of a vertex cover is also a vertex cover. □

We now show a hardness result for ONLINE VERTEX COVER. In our reduction, we make use of the following graph construction. The same construction will also be used later on for other problems. We remark that this graph construction is identical to the one used in [19] for showing lower bounds for ONLINE INDEPENDENT SET in the multi-solution model.

Definition 8 (cf. [19]) For any string $x = x_1 \ldots x_n \in \{0,1\}^n$, define $G_x = (V,E)$ as follows:

$$V = \{v_1, \ldots, v_n\},$$
$$E = \{(v_i, v_j) \colon x_i = 1 \text{ and } i < j\}.$$

Furthermore, let $V_0 = \{v_i \colon x_i = 0\}$ and $V_1 = \{v_i \colon x_i = 1\}$.

For a string $x \in \{0,1\}^n$, the graph G_x from Definition 8 is a *split graph*: The vertex set V can be partitioned into V_0 and V_1 such that V_0 is an independent set of size $|x|_0$ and V_1 is a clique of size $|x|_1$.

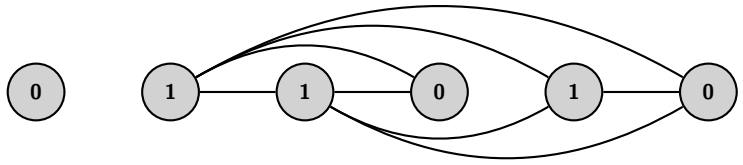

Fig. 2 G_{011010}

Lemma 3 *If there is a c-competitive algorithm reading b bits for* ONLINE VERTEX COVER*, then there is a c-competitive algorithm reading* $b + O(\log n)$ *bits for* MINASGK.

Proof Let ALG be a c-competitive algorithm for ONLINE VERTEX COVER reading at most b bits of advice. By definition, there exists a constant α such that $\text{ALG}(\sigma) \leq c \cdot \text{OPT}(\sigma) + \alpha$ for any input sequence σ. We will define an algorithm, ALG$'$, and an oracle, $0'$, for MINASGK such that ALG$'$ is c-competitive (with the same additive constant) and reads at most $b + O(\log n)$ bits of advice.

For $x = x_1 \ldots x_n$ an input string to MINASGK, consider the input instance to ONLINE VERTEX COVER $G_x = (V,E)$ defined in Definition 8 where the vertices are requested in the order $\langle v_1, \ldots, v_n \rangle$. We say that a vertex in V_0 is *bad* and that a vertex in V_1 is *good*. Note that $V_1 \setminus \{v_n\}$ is a minimum vertex cover of G_x. Also, if an algorithm rejects a good vertex v_i, then it must accept all later vertices v_j (where $i < j \leq n$) in order to cover the edges (v_i, v_j). In particular, since the good vertices form a clique, no algorithm can reject more than one good vertex.

Let φ be the advice read by ALG, and let V_{ALG} be the vertices chosen by ALG. Since ALG is c-competitive, we know that V_{ALG} must be a vertex cover of size at most $c\,|V_1 \setminus \{v_n\}| + \alpha \leq c\,|V_1| + \alpha$.

We now define ALG$'$ and $0'$. As usual, y denotes the output computed by ALG$'$. We consider three cases. The first two bits of the advice tape will be used to tell ALG$'$ which one of the three cases we are in.

Case 1: ALG *accepts all good vertices in* G_x*, i.e.,* $V_1 \subseteq V_{\text{ALG}}$. The oracle $0'$ writes the advice φ to the advice tape. When ALG$'$ receives request i, it considers what ALG does when the vertex v_i in G_x is revealed: ALG$'$ answers 1 if ALG accepts v_i and 0 otherwise. Note that it is possible for ALG$'$ to simulate ALG since, at the beginning of round i, ALG$'$ knows $x_1 \ldots x_{i-1}$. In particular, ALG$'$ knows which edges to reveal to ALG along with the vertex v_i in G_x. Together with access to the advice φ read by ALG, this allows ALG$'$ to simulate ALG. Since $V_1 \subseteq V_{\text{ALG}}$, we get that $x \sqsubseteq y$. Furthermore, since $|V_{\text{ALG}}| \leq c\,|V_1| + \alpha$, we also get that $|y|_1 \leq c\,|x|_1 + \alpha$.

Case 2a: ALG *rejects a good vertex,* v_i*, and accepts a bad vertex,* v_j. In this case, the oracle $0'$ writes the indices of i and j in a self-delimiting way, followed by φ,

to the advice tape. ALG′ simulates ALG as before and answers accordingly, except that it answers 1 in round i and 0 in round j. This ensures that $x \sqsubseteq y$. Furthermore, $|y|_1 = |V_{\mathrm{ALG}}| \leq c\,|V_1| + \alpha = c\,|x|_1 + \alpha$.

Case 2b: ALG *rejects a good vertex, v_i, and all bad vertices.* In this case, $V_{\mathrm{ALG}} = V_1 \setminus \{v_i\}$. The oracle O' writes the value of i to the advice tape in a self-delimiting way, followed by φ. Again, ALG′ simulates ALG, but it answers 1 in round i. Thus, $x = y$, meaning that $x \sqsubseteq y$ and y is optimal.

In all cases, ALG′ computes an output y such that $x \sqsubseteq y$ and $|y|_1 \leq c\,|x|_1 + \alpha$. Since $|\varphi| \leq b$, the maximum number of bits read by ALG′ is $b + O(\log n) + 2 = b + O(\log n)$. □

Theorem 12 ONLINE VERTEX COVER *is* AOC*-complete.*

Proof By Lemma 2, ONLINE VERTEX COVER is in AOC. Combining Lemma 3 and Theorem 5 shows that a c-competitive algorithm for ONLINE VERTEX COVER must read at least $B(n,c) - O(\log n)$ bits of advice. Thus, ONLINE VERTEX COVER is AOC-complete. □

4.1.2 Online Cycle Finding.

Most of the graph problems that we prove to be AOC-complete are, in their offline versions, NP-complete. However, in this section, we show that ONLINE CYCLE FINDING is also AOC-complete. The offline version of this problem is very simple and can easily be solved in polynomial time.

Lemma 4 ONLINE CYCLE FINDING *is in* AOC.

This and the following proofs of membership of AOC have been omitted. They are almost identical to the proof of Lemma 2.

In order to show that ONLINE CYCLE FINDING is AOC-complete, we will make use of the following graph.

Definition 9 For a string $x = x_1 \ldots x_n \in \{0,1\}^n$ define $f(x_i)$ to be the largest $j < i$ such that $x_j = 1$. Note that this may not always be defined. We let MAX be the largest i such that $x_i = 1$. Similarly, we let MIN be the smallest i such that $x_i = 1$. We now define the graph $H_x = (V,E)$:

$$V = \{v_1, \ldots, v_n\},$$
$$E = \{(v_j, v_i) \colon f(x_i) = j\} \cup \{(v_{\mathrm{MIN}}, v_{\mathrm{MAX}})\}.$$

Furthermore, let $V_0 = \{v_i \colon x_i = 0\}$ and $V_1 = \{v_i \colon x_i = 1\}$.

Lemma 5 *If there is a c-competitive algorithm reading b bits for* ONLINE CYCLE FINDING*, then there is a c-competitive algorithm reading $b + O(\log n)$ bits for* MIN-ASGK.

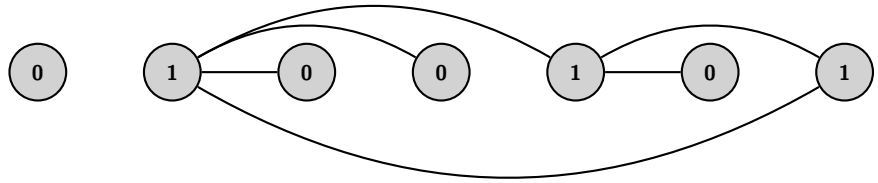

Fig. 3 $H_{0100101}$

Proof Let ALG be a c-competitive algorithm (with an additive constant α) for ONLINE CYCLE FINDING reading at most b bits of advice. We will define an algorithm ALG$'$ and an oracle $0'$ for MINASGK such that ALG$'$ is c-competitive (with the same additive constant) and reads at most $b + O(\log n)$ bits of advice.

Let $x = x_1 \dots x_n$ be an input string to MINASGK. The oracle $0'$ first writes one bit of advice to indicate if $|x|_1 \leq 2$. If this is the case, $0'$ writes (in a self-delimiting way) the index of these at most two 1s to the advice tape. This can be done using $O(\log n)$ bits and clearly allows ALG$'$ to be strictly 1-competitive. In the rest of the proof, we will assume that there are at least three 1s in x.

Consider the input instance to ONLINE CYCLE FINDING, $H_x = (V,E)$, defined in Definition 9, where the vertices are requested in the order $\langle v_1, \dots, v_n \rangle$. Note that the vertices V_1 form the only cycle in H_x. Thus, if an algorithm rejects a vertex from V_1, the subgraph induced by the vertices accepted by the algorithm cannot contain a cycle.

Let φ be the advice read by ALG, and let V_{ALG} be the vertices chosen by ALG, when the n vertices of H_x are revealed. Since ALG is c-competitive, we know that $|V_{\text{ALG}}| \leq c\,|V_1| + \alpha$.

We now define ALG$'$. As usual, y denotes the output computed by ALG$'$. Since ALG is c-competitive, it must hold that $V_1 \subseteq V_{\text{ALG}}$. The oracle $0'$ writes the advice φ to the advice tape. When ALG$'$ receives request i at the beginning of round i in MINASGK, it considers what ALG does when the vertex v_i in H_x is revealed: ALG$'$ answers 1 if ALG accepts v_i and 0 otherwise. Note that it is possible for ALG$'$ to simulate ALG since, at the beginning of round i, ALG$'$ knows $x_1 \dots x_{i-1}$. In particular, ALG$'$ knows which edges were revealed to ALG along with the vertex v_i in H_x. Note, however, that in order to simulate the edge from v_{MIN} to v_{MAX}, ALG needs to know when v_{MAX} is being revealed. This can be achieved using $O(\log n)$ additional advice bits.

Together with access to the advice φ read by ALG, this allows ALG$'$ to simulate ALG. Since $V_1 \subseteq V_{\text{ALG}}$, we get that $x \sqsubseteq y$. Furthermore, since $|V_{\text{ALG}}| \leq c\,|V_1| + \alpha$, we also get that $|y|_1 \leq c\,|x|_1 + \alpha$. □

Theorem 13 ONLINE CYCLE FINDING *is* AOC*-complete*

Proof This follows from Lemmas 4 and 5 together with Theorem 5. □

4.1.3 Online Dominating Set.

In this section, we show that ONLINE DOMINATING SET is also AOC-complete. We do not require that the vertices picked by the online algorithm form a dominating set at all times. We only require that the solution produced by the algorithm is a dominating set when the request sequence ends. Of course, this makes a difference only because we consider online algorithms with advice. For ONLINE VERTEX COVER, this issue did not arise, since it is not possible to end up with a vertex cover without maintaining a vertex cover at all times. Thus, in this aspect, ONLINE DOMINATING SET is more similar to ONLINE CYCLE FINDING.

Lemma 6 ONLINE DOMINATING SET *in in* AOC.

In order to show that ONLINE DOMINATING SET is AOC-complete, we use the following construction.

Definition 10 For a string $x = x_1 \dots x_n$ such that $|x|_1 \geq 1$, define MAX to be the largest i such that $x_i = 1$ and define $K_x = (V,E)$ as follows:

$$V = \{v_1, \dots, v_n\},$$
$$E = \{(v_i, v_{\text{MAX}}) \colon x_i = 0\}.$$

Furthermore, let $V_0 = \{v_i \colon x_i = 0\}$ and $V_1 = \{v_i \colon x_i = 1\}$.

Note that V_1 is a smallest dominating set in K_x and that any dominating set is either a superset of V_1 or equal to $V \setminus \{v_{\text{MAX}}\}$. We now give a lower bound on the advice complexity of ONLINE DOMINATING SET. Interestingly, it is possible to do this by making a reduction from MINASGU (instead of MINASGK) to ONLINE DOMINATING SET.

Lemma 7 *If there is a c-competitive algorithm for* ONLINE DOMINATING SET *reading b bits of advice, then there is a c-competitive algorithm reading $b + O(\log n)$ bits of advice for* MINASGU.

Proof Let ALG be a c-competitive algorithm (with an additive constant of α) for ONLINE DOMINATING SET reading at most b bits of advice. We will define an algorithm ALG$'$ and an oracle $0'$ for MINASGU such that ALG$'$ is c-competitive (with the same additive constant) and reads at most $b + O(\log n)$ bits of advice.

Let $x = x_1 \dots x_n$ be an input string to MINASGU. The oracle $0'$ first writes one bit of advice to indicate if $|x|_1 = 0$. If this is the case, ALG$'$ answers 0 in each round. In the rest of the proof, we will assume that $|x|_1 \geq 1$.

Consider the input instance to ONLINE DOMINATING SET, $K_x = (V,E)$, defined in Definition 10, where the vertices are requested in the order $\langle v_1, \dots, v_n \rangle$. Note that V_1 is the smallest dominating set in K_x. Let φ be the advice read by ALG, and let V_{ALG} be the vertices chosen by ALG, when the n vertices of K_x are revealed. Since ALG is c-competitive, we know that V_{ALG} is a dominating set of size $|V_{\text{ALG}}| \leq c\,|V_1| + \alpha$.

We now define ALG$'$ and $0'$. The second bit of the advice tape will be used to let ALG$'$ distinguish the two cases described below. Note that the only vertex from V_1

that can be rejected by a c-competitive algorithm is v_{MAX}, and nothing can be rejected when $V_1 = V$. Hence the two cases are exhaustive.

Case 1: ALG *accepts all vertices in* V_1. The oracle $\mathtt{O}'$ writes the value of MAX in a self-delimiting way. This requires $O(\log n)$ bits. Furthermore, $\mathtt{O}'$ writes φ to the advice tape. Now, ALG$'$ learns φ and MAX and works as follows: In round $i \leq \text{MAX} - 1$, ALG$'$ answers 1 if ALG accepts the vertex v_i and 0 otherwise. Note that ALG$'$ knows that no edges are revealed to ALG in the first MAX $-$ 1 rounds. Thus, ALG$'$ can compute the answer produced by ALG in these rounds from φ alone. In round MAX, ALG$'$ answers 1. In rounds $\text{MAX}+1,\ldots,n$, the algorithm ALG$'$ always answers 0.

Case 2: ALG *rejects* v_{MAX}. In order to dominate v_{MAX}, ALG must accept a vertex $v_i \in V_0$. The oracle $\mathtt{O}'$ writes the values of MAX and i in a self-delimiting way, followed by φ, to the advice tape. ALG$'$ behaves as in Case 1, except that it answers 0 in round i.

In both cases, $x \sqsubseteq y$ and $|y|_1 \leq |V_{\text{ALG}}| \leq c\,|V_1| + \alpha = c\,|x|_1 + \alpha$. Furthermore, ALG$'$ reads $b + O(\log n)$ bits of advice. □

Theorem 14 ONLINE DOMINATING SET *is* AOC*-complete*

Proof This follows from Lemmas 6 and 7 together with Theorem 5. □

4.1.4 Online Set Cover.

We study a version of ONLINE SET COVER in which the universe is known from the beginning and the sets arrive online. Note that this problem is very different from the set cover problem studied in [1, 24], where the elements (and not the sets) arrive online.

Lemma 8 ONLINE SET COVER *is in* AOC

Lemma 9 *If there is a c-competitive algorithm for* ONLINE SET COVER *reading b bits of advice, then there is a c-competitive algorithm reading $b + O(\log n)$ bits of advice for* MINASGU.

Proof Let $x = x_1 \ldots x_n$ be an input string to MINASGU with $|x|_1 \geq 1$, and define MAX as in Definition 10. We define an instance of ONLINE SET COVER as follows. The universe is $[n] = \{1,\ldots,n\}$ and there are n requests. For $i \neq \text{MAX}$, request i is just the singleton $\{i\}$. Request MAX is the set $\{\text{MAX}\} \cup S_0$, where $S_0 = \{i : x_i = 0\}$.

Using these instances of ONLINE SET COVER and the same arguments as in Lemma 7 proves the theorem. Note that for ONLINE SET COVER, only Case 1 of Lemma 7 is relevant, since a c-competitive algorithm for this problem will accept all requests $i \notin S_0$. □

Theorem 15 ONLINE SET COVER *is* AOC*-complete*

Proof This follows from Lemmas 8 and 9 together with Theorem 5. □

4.2 AOC-complete maximization problems

In this section, we consider two maximization problems which are AOC-complete.

4.2.1 Online Independent Set.

The first maximization problem that we consider is ONLINE INDEPENDENT SET.

Lemma 10 ONLINE INDEPENDENT SET *is in* AOC.

Proof Each request is a vertex along with the edges connecting it to previously requested vertices. The algorithm makes a binary choice for each request, to include the vertex ($y_i = 0$) or not ($y_i = 1$). The feasible outputs are those that are independent sets. There exists a feasible output (taking no vertices). The score of a feasible output is the number of vertices in it, and the score of an infeasible output is $-\infty$. Any subset of the vertices in an optimal solution is a feasible solution. □

Lemma 11 *If there is a c-competitive algorithm reading b bits for* ONLINE INDEPENDENT SET*, then there is a c-competitive algorithm reading* $b + O(\log n)$ *bits for* MAXASGK.

Proof The proof is almost identical to the proof of Lemma 3. Let ALG be a c-competitive algorithm (with an additive constant of α) for ONLINE INDEPENDENT SET reading at most b bits of advice. We will define an algorithm ALG$'$ and an oracle $\mathtt{O}'$ for MAXASGK such that ALG$'$ is c-competitive (with the same additive constant) and reads at most b bits of advice.

As in Lemma 3, on input $x = x_1 \dots x_n$ to ONLINE INDEPENDENT SET, the algorithm ALG$'$ simulates ALG on G_x (from Definition 8). This time, a vertex in V_0 is *good* and a vertex in V_1 is *bad*. Note that $V_0 \cup \{v_n\}$ is a maximum independent set in G_x. Also, if ALG accepts a bad vertex, v_i, then no further vertices v_j (where $i < j \leq n$) can be accepted because of the edges (v_i, v_j). Thus, ALG accepts at most one bad vertex. Let V_{ALG} be the vertices accepted by ALG. Since ALG is c-competitive, V_{ALG} is an independent set satisfying $|V_0| \leq |V_0 \cup \{v_n\}| \leq c\,|V_{\text{ALG}}| + \alpha$. We denote by y the output computed by ALG$'$. There are three cases to consider:

Case 1: All vertices accepted by ALG *are good, that is* $V_{\text{ALG}} \subseteq V_0$. In this case, ALG answers 0 in round i if $v_i \in V_{\text{ALG}}$ and 1 otherwise. Clearly, $x \sqsubseteq y$ and $|x|_0 = |V_0| \leq c\,|V_{\text{ALG}}| + \alpha = c\,|y|_0 + \alpha$.

Case 2a: ALG *accepts a bad vertex,* v_i*, and rejects a good vertex,* v_j. In this case, the oracle $\mathtt{O}'$ writes the indices i and j in a self-delimiting way. ALG simulates ALG$'$ as before, but answers 1 in round i and 0 in round j. It follows that $x \sqsubseteq y$ and $|x|_0 \leq c\,|y|_0 + \alpha$.

Case 2b: ALG *accepts a bad vertex,* v_i*, and all good vertices.* This implies that V_{ALG} is an independent set of size $|V_0| + 1$, which must be optimal. The oracle $\mathtt{O}'$ writes the value of i to the advice tape in a self-delimiting way. ALG simulates ALG$'$ as before but answers 1 in round i. It follows that $x \sqsubseteq y$. Furthermore, $|y|_0 = |V_{\text{ALG}}| - 1 = |V_0| = |x|_0$, and hence the solution y is optimal.

$$x = 010$$
$$L = 8$$
$$I_x = \langle (0,4),(4,6),(4,5) \rangle$$

Fig. 4 An example of the reduction used in the proof of Lemma 13. The request $(0,4)$ is good, since $x_1 = 0$, and $(4,6)$ is a bad request, since $x_2 = 1$.

In order to simulate ALG, the algorithm ALG$'$ needs to read at most b bits of advice plus $O(\log n)$ bits of advice to specify the case and handle the cases where ALG accepts a bad vertex. □

Theorem 16 ONLINE INDEPENDENT SET *is* AOC*-complete.*

Proof This follows from Lemmas 10 and 11 together with Theorem 8. □

4.2.2 Online Disjoint Path Allocation.

In this section, we show that ONLINE DISJOINT PATH ALLOCATION is AOC-complete.

Lemma 12 ONLINE DISJOINT PATH ALLOCATION *is in* AOC.

In Lemma 13, we use the same hard instance for ONLINE DISJOINT PATH ALLOCATION as in [7] to get a lower bound on the advice complexity of ONLINE DISJOINT PATH ALLOCATION.

Lemma 13 *If there is a c-competitive algorithm reading b bits for* ONLINE DISJOINT PATH ALLOCATION*, then there is a c-competitive algorithm reading $b + O(\log n)$ bits for* MAXASGK.

Proof The proof is similar to the proof of Lemma 11. Let ALG be a c-competitive algorithm for ONLINE DISJOINT PATH ALLOCATION reading at most b bits of advice. We will describe an algorithm ALG$'$ and an oracle $0'$ for MAXASGK such that ALG$'$ is c-competitive (with the same additive constant α as ALG) and reads at most b bits of advice.

Let $x = x_1 \dots x_n$ be an input to MAXASGK. We define an instance I_x of ONLINE DISJOINT PATH ALLOCATION with $L = 2^n$ (that is, the number of vertices on the path is $2^n + 1$). In round i, $1 \leq i \leq n$, a path of length 2^{n-i} arrives. For $2 \leq i \leq n$, the position of the path depends on x_{i-1}. We define the request sequence inductively (for an example, see Figure 4):

In round 1, the request (u_1, v_1) arrives, where

$$u_1 = 0$$
$$v_1 = 2^{n-1}$$

In round i, $2 \leq i \leq n$, the request (u_i, v_i) arrives, where

$$u_i = \begin{cases} u_{i-1}, & \text{if } x_{i-1} = 1 \\ v_{i-1}, & \text{if } x_{i-1} = 0 \end{cases}$$
$$v_i = u_i + 2^{n-i}$$

We say that a request $r_i = (u_i, v_i)$ is *good* if $x_i = 0$ and *bad* if $x_i = 1$. If r_i is good, then none of the later requests overlap with r_i. On the other hand, if r_i is bad, then all later requests do overlap with r_i. In particular, if one accepts a single bad request, then no further requests can be accepted. An optimal solution is obtained if one accepts all good requests together with r_n.

The oracle $\mathtt{O}'$ will provide ALG$'$ with the advice φ read by ALG when processing I_x. Since ALG knows the value of x_{i-1} at the beginning of round i in MAXASGK, ALG can use the advice φ to simulate ALG$'$ on I_x. If ALG only accepts good requests, it is clear that ALG$'$ can compute an output y such that $x \sqsubseteq y$ and $|x|_0 \leq c\,|y|_0 + \alpha$. The case where ALG accepts a bad request is handled by using at most $O(\log n)$ additional advice bits, exactly as in Lemma 11. □

Theorem 17 ONLINE DISJOINT PATH ALLOCATION *is* AOC*-complete.*

4.3 AOC Problems which are not AOC-complete

In this section, we will give two examples of problems in AOC which are provably not AOC-complete.

4.3.1 Uniform knapsack.

We define the problem ONLINE UNIFORM KNAPSACK as follows: For each request, i, an item of weight a_i, $0 \leq a_i \leq 1$, is requested. A request must immediately be either accepted or rejected, and this decision is irrevocable. Let S denote the set of indices of accepted items. We say that S is a feasible solution if $\sum_{i \in S} a_i \leq 1$. The profit of a feasible solution is the number of items accepted (all items have a value of 1). The problem is a maximization problem.

The ONLINE UNIFORM KNAPSACK problem is the online knapsack problem as studied in [8], but with the restriction that all items have a value of 1. This problem is the same as online dual bin packing with only a single bin available and where items can be rejected.

It is clear that ONLINE UNIFORM KNAPSACK belongs to AOC since a subset of a feasible solution is also a feasible solution. Furthermore, since all items have value 1, the profit of a feasible solution is simply the number of items packed in the knapsack. The problem is hard in the sense that no deterministic algorithm (without advice) can attain a strict competitive ratio better than $\Omega(n)$ (see [8, 25]). However, as the next lemma shows, the problem is not AOC-complete. In [8], it is shown that for any $\varepsilon > 0$, it is possible to achieve a competitive ratio of $1 + \varepsilon$ using $O(\log n)$ bits of advice, under the assumption that all weights and values can be represented

in polynomial space. Lemma 14 shows how this assumption can be avoided when all items have unit value.

Lemma 14 *There is a strictly* 2*-competitive* ONLINE UNIFORM KNAPSACK *algorithm reading* $O(\log n)$ *bits of advice, where* n *is the length of the input.*

Proof Fix an input $\sigma = \langle a_1, \ldots, a_n \rangle$. Let m be the number of items accepted by OPT. The oracle writes m to the advice tape using a self-delimiting encoding. Since $m \leq n$, this requires $O(\log n)$ bits. The algorithm ALG learns m from the advice tape and works as follows: If ALG is offered an item, a_i, such that $a_i \leq 2/m$ and if accepting a_i will not make the total weight of ALG's solution larger than 1, then ALG accepts a_i. Otherwise, a_i is rejected.

In order to show that ALG is strictly 2-competitive, we define $A = \{a_i : a_i \leq 2/m\}$. First note that $|A| \geq m/2$, since the sizes of the m smallest items add up to at most 1. Thus, if ALG accepts all items contained in A, it accepts at least $m/2$ items. On the other hand, if ALG rejects any item $a_i \in A$, it means that it has already accepted items of total size more than $1 - 2/m$. Since all accepted items have size at most $2/m$, this means that ALG has accepted at least $m/2$ items. □

Even though ONLINE UNIFORM KNAPSACK is not AOC-complete, the fact that it belongs to AOC might still be of interest, since this provides some starting point for determining the advice complexity of the problem. In particular, it gives some (non-trivial) way to obtain a c-competitive algorithm for $c < 2$. Determining the exact advice complexity of ONLINE UNIFORM KNAPSACK is left as an open problem.

4.3.2 Matching under edge-arrival.

We briefly consider the ONLINE MATCHING problem in an edge-arrival version. For each request, an edge is revealed. An edge can be either accepted or rejected. Denote by E_{ALG} the edges accepted by some algorithm ALG. A solution E_{ALG} is feasible if the set of edges in the solution is a matching in the input graph, and the profit of a feasible solution is the number of edges in E_{ALG}. The problem is a maximization problem.

It is well-known that the greedy algorithm is 2-competitive for ONLINE MATCHING. Since this algorithm works in an online setting without any advice, it follows that ONLINE MATCHING is not AOC-complete. On the other hand, ONLINE MATCHING is in AOC. This gives an upper bound on the advice complexity of the problem for $1 \leq c < 2$. It seems obvious that this upper bound is not tight, but currently, no better bound is known.

5 Conclusion and Open Problems

The following theorem summarizes the main results of this paper.

Theorem 18 *For the problems*

- ONLINE VERTEX COVER

- ONLINE CYCLE FINDING
- ONLINE DOMINATING SET
- ONLINE SET COVER *(set-arrival version)*
- ONLINE INDEPENDENT SET
- ONLINE DISJOINT PATH ALLOCATION

and for any $c > 1$, possibly a function of the input length n,

$$b = \log\left(1 + \frac{(c-1)^{c-1}}{c^c}\right) n \pm O(\log n)$$

bits of advice are necessary and sufficient to achieve a (strict) competitive ratio of c.

As with the original string guessing problem SG [5,15], we have shown that ASG is a useful tool for determining the advice complexity of online problems. It seems plausible that one could identify other variants of online string guessing and obtain classes similar to AOC. Potentially, this could lead to an entire hierarchy of string guessing problems and related classes.

More concretely, there are various possibilities of generalizing ASG. One could associate some positive weight to each bit x_i in the input string. The goal would then be to produce a feasible output of minimum (or maximum) weight. Such a string guessing problem would model minimum weight vertex cover (or maximum weight independent set). Note that for MAXASG, the algorithm from Theorem 2 works in the weighted version. However, the same is not true for any of the algorithms we have given for MINASG. Thus, it remains an open problem if $O(n/c)$ bits of advice suffice to achieve a competitive ratio of c for the weighted version of MINASG.

Acknowledgements The authors would like to thank Magnus Gausdal Find for helpful discussions.

References

1. Alon, N., Awerbuch, B., Azar, Y., Buchbinder, N., Naor, J.: The online set cover problem. SIAM J. Comput. **39**(2), 361–370 (2009)
2. Barhum, K.: Tight bounds for the advice complexity of the online minimum steiner tree problem. In: Proc. 40th International Conf. on Current Trends in Theory and Practice of Computer Science (SOFSEM), *Lecture Notes in Comput. Sci., Springer*, vol. 8327, pp. 77–88 (2014)
3. Barhum, K., Böckenhauer, H.J., Forišek, M., Gebauer, H., Hromkovič, J., Krug, S., Smula, J., Steffen, B.: On the power of advice and randomization for the disjoint path allocation problem. In: Proc. 40th International Conf. on Current Trends in Theory and Practice of Computer Science (SOFSEM), *Lecture Notes in Comput. Sci., Springer*, vol. 8327, pp. 89–101 (2014)
4. Bianchi, M.P., Böckenhauer, H.J., Hromkovič, J., Keller, L.: Online coloring of bipartite graphs with and without advice. Algorithmica **70**(1), 92–111 (2014)
5. Böckenhauer, H.J., Hromkovič, J., Komm, D., Krug, S., Smula, J., Sprock, A.: The string guessing problem as a method to prove lower bounds on the advice complexity. Theor. Comput. Sci. **554**, 95–108 (2014)
6. Böckenhauer, H.J., Komm, D., Královič, R., Královič, R.: On the advice complexity of the k-server problem. In: Proc. 38th International Colloquium on Automata, Languages, and Programming (ICALP), *Lecture Notes in Comput. Sci., Springer*, vol. 6755, pp. 207–218 (2011)
7. Böckenhauer, H.J., Komm, D., Královič, R., Královič, R., Mömke, T.: On the advice complexity of online problems. In: Proc. 20th International Symp. on Algorithms and Computation (ISAAC), *Lecture Notes in Comput. Sci., Springer*, vol. 5878, pp. 331–340 (2009)

8. Böckenhauer, H.J., Komm, D., Královič, R., Rossmanith, P.: The online knapsack problem: Advice and randomization. Theor. Comput. Sci. **527**, 61–72 (2014)
9. Boyar, J., Kamali, S., Larsen, K.S., López-Ortiz, A.: On the list update problem with advice. In: Proc. 8th International Conf. on Language and Automata Theory and Applications (LATA), *Lecture Notes in Comput. Sci., Springer*, vol. 8370, pp. 210–221 (2014). Full paper to appear in *Information and Computation*.
10. Boyar, J., Kamali, S., Larsen, K.S., López-Ortiz, A.: Online bin packing with advice. Algorithmica **74**, 507–527 (2016)
11. Demange, M., Paschos, V.T.: On-line vertex-covering. Theor. Comput. Sci. **332**(1-3), 83–108 (2005)
12. Dinitz, J.H., Stinson, D.R. (eds.): Contemporary Design Theory: a Collection of Surveys. Wiley-Interscience series in discrete mathematics and optimization. Wiley, New York (1992). URL `http://opac.inria.fr/record=b1088981`
13. Dobrev, S., Královič, R., Královič, R.: Advice complexity of maximum independent set in sparse and bipartite graphs. Theory Comput. Syst. **56**(1), 197–219 (2015)
14. Dobrev, S., Královič, R., Pardubská, D.: Measuring the problem-relevant information in input. RAIRO - Theor. Inf. Appl. **43**(3), 585–613 (2009)
15. Emek, Y., Fraigniaud, P., Korman, A., Rosén, A.: Online computation with advice. Theor. Comput. Sci. **412**(24), 2642–2656 (2011)
16. Erdős, P., Spencer, J.: Probabilistic Methods in Combinatorics. Academic Press (1974)
17. Forišek, M., Keller, L., Steinová, M.: Advice complexity of online coloring for paths. In: Proc. 6th International Conf. on Language and Automata Theory and Applications (LATA), *Lecture Notes in Comput. Sci., Springer*, vol. 7183, pp. 228–239 (2012)
18. Gupta, S., Kamali, S., López-Ortiz, A.: On advice complexity of the k-server problem under sparse metrics. In: Proc. 20th International Colloquium on Structural Information and Communication Complexity (SIROCCO), *Lecture Notes in Comput. Sci., Springer*, vol. 8179, pp. 55–67 (2013)
19. Halldórsson, M.M., Iwama, K., Miyazaki, S., Taketomi, S.: Online independent sets. Theor. Comput. Sci. **289**(2), 953–962 (2002)
20. Halldórsson, M.M., Szegedy, M.: Lower bounds for on-line graph coloring. Theor. Comput. Sci. **130**(1), 163–174 (1994)
21. Håstad, J.: Clique is hard to approximate within $n^{1-\varepsilon}$. Acta Math. **182**(1), 105–142 (1999)
22. Hromkovič, J., Královič, R., Královič, R.: Information complexity of online problems. In: Proc. 35th Symp. on Mathematical Foundations of Computer Science (MFCS), *Lecture Notes in Comput. Sci., Springer*, vol. 6281, pp. 24–36 (2010)
23. Karlin, A.R., Manasse, M.S., Rudolph, L., Sleator, D.D.: Competitive snoopy caching. Algorithmica **3**, 77–119 (1988)
24. Komm, D., Královič, R., Mömke, T.: On the advice complexity of the set cover problem. In: Proc. 7th International Computer Science Symp. in Russia (CSR), *Lecture Notes in Comput. Sci., Springer*, vol. 7353, pp. 241–252 (2012)
25. Marchetti-Spaccamela, A., Vercellis, C.: Stochastic on-line knapsack problems. Math. Program. **68**, 73–104 (1995)
26. Mikkelsen, J.W.: Optimal online edge coloring of planar graphs with advice. In: Proc. 9th International Conf. on Algorithms and Complexity (CIAC), *Lecture Notes in Comput. Sci., Springer*, vol. 9079, pp. 352–364 (2015)
27. Mitzenmacher, M., Upfal, E.: Probability and Computing - Randomized Algorithms and Probabilistic Analysis. Cambridge University Press (2005)
28. Miyazaki, S.: On the advice complexity of online bipartite matching and online stable marriage. Inf. Process. Lett. **114**(12), 714–717 (2014)
29. Raz, R., Safra, S.: A sub-constant error-probability low-degree test, and a sub-constant error-probability PCP characterization of NP. In: Proc. 29th Symp. on Theory of Computing (STOC), pp. 475–484. ACM (1997)
30. Renault, M.P., Rosén, A., van Stee, R.: Online algorithms with advice for bin packing and scheduling problems. Theor. Comput. Sci. **600**, 155–170 (2015)
31. Seibert, S., Sprock, A., Unger, W.: Advice complexity of the online coloring problem. In: Proc. 8th International Conf. on Algorithms and Complexity (CIAC), *Lecture Notes in Comput. Sci., Springer*, vol. 7878, pp. 345–357 (2013)
32. Sleator, D.D., Tarjan, R.E.: Amortized efficiency of list update and paging rules. Commun. ACM **28**(2), 202–208 (1985)

Appendix

A Approximation of the Advice Complexity Bounds

In Theorems 3-8, bounds on the advice complexity of ASG were obtained. These bounds are tight up to an additive term of $O(\log n)$. However, within the proofs, they are all expressed in terms of the minimum size of a certain covering design or a quotient of binomial coefficients. In this appendix, we prove the closed formula estimates for the advice complexity stated in Theorems 3-8 and 11. Again, these estimates are tight up to an additive term of $O(\log n)$. The key to obtaining the estimates is the estimation of a binomial coefficient using the binary entropy function.

A.1 Approximating the Function $B(n,c)$

Lemma 15 *For $c > 1$, it holds that*

$$\frac{1}{e\ln(2)}\frac{1}{c} \leq \log\left(1+\frac{(c-1)^{c-1}}{c^c}\right) \leq \frac{1}{c}.$$

Proof We prove the upper bound first. To this end, note that

$$\log\left(1+\frac{(c-1)^{c-1}}{c^c}\right) \leq \frac{1}{c} \Leftrightarrow 1+\frac{(c-1)^{c-1}}{c^c} \leq 2^{1/c} \Leftrightarrow \left(1+\frac{(c-1)^{c-1}}{c^c}\right)^c \leq 2.$$

Using calculus, one may verify that $\left(1+\frac{(c-1)^{c-1}}{c^c}\right)^c$ is decreasing in c for $c > 1$. Thus, by continuity, it follows that

$$\begin{aligned}\left(1+\frac{(c-1)^{c-1}}{c^c}\right)^c &\leq \lim_{c\to 1^+}\left(1+\frac{(c-1)^{c-1}}{c^c}\right)^c = \lim_{c\to 1^+}\left(1+\left(\frac{c-1}{c}\right)^{c-1}\frac{1}{c}\right)^c \\ &= \lim_{c\to 1^+}\left(1+\frac{1}{c}\right)^c = 2.\end{aligned}$$

For the lower bound, let $a = e\ln(2)$ and note that

$$\frac{1}{ac} \leq \log\left(1+\frac{(c-1)^{c-1}}{c^c}\right) \Leftrightarrow 2 \leq \left(1+\frac{(c-1)^{c-1}}{c^c}\right)^{ac}.$$

Again, using calculus, one may verify that $\left(1+\frac{(c-1)^{c-1}}{c^c}\right)^{ac}$ is decreasing in c for $c > 1$. It follows that

$$\begin{aligned}\left(1+\frac{(c-1)^{c-1}}{c^c}\right)^{ac} &\geq \lim_{c\to\infty}\left(1+\frac{(c-1)^{c-1}}{c^c}\right)^{ac} = \lim_{c\to\infty}\left(1+\left(\frac{c-1}{c}\right)^{c-1}\frac{1}{c}\right)^{ac} \\ &= \lim_{c\to\infty}\left(1+\frac{1}{e}\frac{1}{c}\right)^{ac} = \lim_{c\to\infty}\left(1+\frac{a/e}{ac}\right)^{ac} = e^{a/e} = e^{\ln(2)} = 2.\end{aligned}$$

□

A.2 The Binary Entropy Function

In this section, we give some properties of the binary entropy function that will be used extensively in Section A.4.

Definition 11 The *binary entropy function* $H : [0,1] \to [0,1]$ is the function given by

$$H(p) = -p\log(p) - (1-p)\log(1-p), \text{ for } 0 < p < 1,$$

and $H(0) = H(1) = 0$.

Lemma 16 (Lemma 9.2 in [27]) *For integers m,n such that $0 \leq m \leq n$,*

$$\frac{2^{nH(m/n)}}{n+1} \leq \binom{n}{m} \leq 2^{nH(m/n)}.$$

Proposition 1 *The binary entropy function $H(p)$ has the following properties.*

(H1) $H\left(\frac{1}{s}\right) = \log(s) + \frac{1-s}{s}\log(s-1)$ *for* $s > 1$.
(H2) $sH\left(\frac{1}{s}\right) \leq \log s + 2$ *for* $s > 1$.
(H3) $H'(p) = \log\left(\frac{1}{p} - 1\right)$ *and* $H''(p) < 0$ *for* $0 < p < 1$.
(H4) For any fixed $t > 0$, $sH\left(\frac{t}{s}\right)$ is increasing in s for $s > t$.
(H5) $nH\left(\frac{1}{x}\right) - nH\left(\frac{1}{x} + \frac{1}{n}\right) < 3$ *if* $n \geq 3$ *and* $x > 2$.

Proof (H1): Follows from the definition.
(H2): For $s > 1$,

$$\begin{aligned} sH\left(\frac{1}{s}\right) &= s\left(\log s + \frac{1-s}{s}\log(s-1)\right), \text{ by } (H1) \\ &= \log\left(\left(1 + \frac{1}{s-1}\right)^{s-1} s\right) \leq \log(e \cdot s) = \log(e) + \log(s) \leq \log s + 2. \end{aligned}$$

(H3): Note that H is smooth for $0 < p < 1$. The derivative $H'(p)$ can be calculated from the definition. The second-order derivative is

$$H''(p) = \frac{-1}{(1-p)p\ln(2)},$$

which is strictly less than zero for all $0 < p < 1$.
(H4): Fix $t > 0$. The claim follows by showing that the partial derivative of $sH(\frac{t}{s})$ with respect to s is positive for all $s > t$.

$$\begin{aligned} \frac{d}{ds}\left(sH\left(\frac{t}{s}\right)\right) &= H\left(\frac{t}{s}\right) + sH'\left(\frac{t}{s}\right)\left(-\frac{t}{s^2}\right) = H\left(\frac{t}{s}\right) - \frac{t}{s}H'\left(\frac{t}{s}\right) \\ &= -\frac{t}{s}\log\left(\frac{t}{s}\right) - \left(1 - \frac{t}{s}\right)\log\left(1 - \frac{t}{s}\right) - \frac{t}{s}\log\left(\frac{s}{t} - 1\right), \text{ by Def. 11 and } (H3) \\ &= -\log\left(1 - \frac{t}{s}\right) > 0. \end{aligned}$$

(H5): $H(p)$ is increasing for $0 \le p \le \frac{1}{2}$ and decreasing for $\frac{1}{2} \le p \le 1$. If $\frac{1}{x}+\frac{1}{n} \le \frac{1}{2}$, then the claim is trivially true (since then the difference is negative). Assume therefore that $\frac{1}{x}+\frac{1}{n} > \frac{1}{2}$. Under this assumption, $H(\frac{1}{x})$ increases and $H(\frac{1}{x}+\frac{1}{n})$ decreases as x tends to 2. Thus, $H(\frac{1}{x}) - H(\frac{1}{x}+\frac{1}{n})$ increases as x tends to 2 and, hence,

$$H\left(\frac{1}{x}\right) - H\left(\frac{1}{x}+\frac{1}{n}\right) \le H\left(\frac{1}{2}\right) - H\left(\frac{1}{2}+\frac{1}{n}\right). \tag{3}$$

Inserting into the definition of H gives

$$\begin{aligned} H\left(\frac{1}{2}\right) - H\left(\frac{1}{2}+\frac{1}{n}\right) &= 1 - \left(-\left(\frac{1}{2}+\frac{1}{n}\right)\log\left(\frac{1}{2}+\frac{1}{n}\right) - \left(\frac{1}{2}-\frac{1}{n}\right)\log\left(\frac{1}{2}-\frac{1}{n}\right)\right) \\ &= \frac{1}{n}\log\left(\frac{\frac{1}{2}+\frac{1}{n}}{\frac{1}{2}-\frac{1}{n}}\right) + \frac{1}{2}\log\left(\left(\frac{1}{2}+\frac{1}{n}\right)\left(\frac{1}{2}-\frac{1}{n}\right)\right) + 1 \\ &= \frac{1}{n}\log\left(\frac{n+2}{n-2}\right) + \frac{1}{2}\log\left(\frac{n^2-4}{4n^2}\right) + 1 \end{aligned}$$

Since $(n+2)/(n-2)$ is decreasing for $n \ge 3$, it follows that $\log((n+2)/(n-2)) \le \log(5)$. Furthermore, $(n^2-4)/(4n^2) \le \frac{1}{4}$ for all $n \ge 3$, and so $\frac{1}{2}\log\left((n^2-4)/(4n^2)\right) + 1 \le 0$. We conclude that, for all $n \ge 3$,

$$H\left(\frac{1}{2}\right) - H\left(\frac{1}{2}+\frac{1}{n}\right) \le \frac{\log(5)}{n} < \frac{3}{n}. \tag{4}$$

Combining (3) and (4) proves (H5). □

A.3 Binomial Coefficients

The following proposition is a collection of simple facts about the binomial coefficient that will be used in Sections A.4 and A.5.

Proposition 2 *Let $a, b, c \in \mathbb{N}$.*

(B1) $\binom{a}{b} = \frac{a}{a-b}\binom{a-1}{b}$, *where $b < a$.*
(B2) For fixed b, $\binom{a}{b}$ is increasing in a.
(B3) If $c \le b \le a$, then

$$\frac{\binom{a}{c}}{\binom{b}{c}} = \frac{\binom{a}{b}}{\binom{a-c}{a-b}}.$$

Proof First, we prove $(B1)$:

$$\binom{a}{b} = \frac{a!}{b!(a-b)!} = \frac{a}{a-b}\,\frac{(a-1)!}{b!(a-1-b)!} = \frac{a}{a-b}\binom{a-1}{b}$$

$(B2)$ follows directly from $(B1)$.

To prove $(B3)$, we calculate the two fractions separately:

$$\frac{\binom{a}{c}}{\binom{b}{c}} = \frac{a!}{c!(a-c)!}\frac{c!(b-c)!}{b!} = \frac{a!}{(a-c)!}\frac{(b-c)!}{b!}$$
$$\frac{\binom{a}{b}}{\binom{a-c}{a-b}} = \frac{a!}{b!(a-b)!}\frac{(a-b)!(b-c)!}{(a-c)!} = \frac{a!}{b!}\frac{(b-c)!}{(a-c)!} = \frac{\binom{a}{c}}{\binom{b}{c}}$$

□

A.4 Approximating the Advice Complexity Bounds for MINASG

The following lemma is used for proving Theorems 3–5.

Lemma 17 *For $c > 1$ and $n \geq 3$,*

$$\log\left(\max_{t\colon \lfloor ct\rfloor < n} C(n,\lfloor ct\rfloor,t)\right) \geq \log\left(\max_{t\colon \lfloor ct\rfloor < n} \frac{\binom{n}{t}}{\binom{\lfloor ct\rfloor}{t}}\right) \tag{5}$$
$$\geq \log\left(1+\frac{(c-1)^{c-1}}{c^c}\right)n - 2\log(n+1) - 5 \tag{6}$$

and

$$\log\left(\max_{t\colon \lfloor ct\rfloor < n} C(n,\lfloor ct\rfloor,t)\right) \leq \log\left(\max_{t\colon \lfloor ct\rfloor < n} \frac{\binom{n}{t}}{\binom{\lfloor ct\rfloor}{t}}n\right) \tag{7}$$
$$\leq \log\left(1+\frac{(c-1)^{c-1}}{c^c}\right)n + 3\log(n+1). \tag{8}$$

Proof We prove the upper and lower bounds separately.

Upper bound: Fix n, c. By Lemma 1,

$$C(n,\lfloor ct\rfloor,t) \leq \frac{\binom{n}{t}}{\binom{\lfloor ct\rfloor}{t}}\left(1+\ln\binom{\lfloor ct\rfloor}{t}\right).$$

Note that $1+\ln\binom{\lfloor ct\rfloor}{t} \leq n$ since we consider only $\lfloor ct\rfloor < n$. This proves (7). Now, taking the logarithm on both sides gives

$$\begin{aligned}
\log(C(n,\lfloor ct\rfloor,t)) &\leq \log\left(\frac{\binom{n}{t}}{\binom{\lfloor ct\rfloor}{t}}\right) + \log n \leq \log\left(\frac{\binom{n}{t}}{\binom{\lceil ct\rceil-1}{t}}\right) + \log n \\
&\leq \log\left(\frac{\binom{n}{t}}{\frac{\lceil ct\rceil - t}{\lceil ct\rceil}\binom{\lceil ct\rceil}{t}}\right) + \log n, \text{ by } (B1) \\
&\leq \log\left(\frac{\binom{n}{t}}{\binom{\lceil ct\rceil}{t}}\right) + \log\left(\frac{\lceil ct\rceil}{\lceil ct\rceil - t}\right) + \log n \\
&\leq \log\left(\frac{\binom{n}{t}}{\binom{\lceil ct\rceil}{t}}\right) + 2\log n. \qquad (9)
\end{aligned}$$

Above, we have increased $\lfloor ct \rfloor$ to $\lceil ct \rceil$ in the binomial coefficient (at the price of an additive term of $\log n$). This is done since it will later be convenient to use that $ct \leq \lceil ct \rceil$. Using Lemma 16, we get that

$$\frac{\binom{n}{t}}{\binom{\lceil ct \rceil}{t}} \leq \frac{2^{nH(t/n)}}{2^{\lceil ct \rceil H(t/\lceil ct \rceil)}} \left(\lceil ct \rceil + 1 \right),$$

and therefore

$$\begin{aligned} \log \left(\frac{\binom{n}{t}}{\binom{\lceil ct \rceil}{t}} \right) &\leq nH\left(\frac{t}{n}\right) - \lceil ct \rceil H\left(\frac{t}{\lceil ct \rceil}\right) + \log\left(\lceil ct \rceil + 1\right) \\ &\leq nH\left(\frac{t}{n}\right) - ctH\left(\frac{1}{c}\right) + \log(n+1), \text{ by } \textit{(H4)}. \end{aligned} \tag{10}$$

Define

$$M(n,t) = nH\left(\frac{t}{n}\right) - ctH\left(\frac{1}{c}\right).$$

Combining (9) and (10) shows that

$$\log(C(n, \lfloor ct \rfloor, t)) \leq M(n,t) + 3\log(n+1). \tag{11}$$

The function M is smooth. For any given input length n, we can determine the value of t maximizing $M(n,t)$ using calculus. In order to simplify the notation for these calculations, define

$$x = \left(\frac{c}{c-1}\right)^c (c-1) + 1,$$

and note that

$$\begin{aligned} \log(x-1) &= c\left(\log c + \frac{1-c}{c}\log(c-1)\right) \\ &= cH\left(\frac{1}{c}\right), \text{ by } \textit{(H1)}. \end{aligned} \tag{12}$$

We want to determine those values of t for which $\frac{d}{dt}M(n,t) = 0$:

$$\begin{aligned} & \frac{d}{dt}M(n,t) = \frac{d}{dt}\left(nH\left(\frac{t}{n}\right) - ctH\left(\frac{1}{c}\right)\right) = 0 \\ \Leftrightarrow\ & nH'\left(\frac{t}{n}\right)\cdot\frac{1}{n} - cH\left(\frac{1}{c}\right) = 0 \\ \Leftrightarrow\ & \log\left(\frac{n}{t} - 1\right) = cH\left(\frac{1}{c}\right), \text{ by } \textit{(H3)} \\ \Leftrightarrow\ & \frac{n}{t} = 2^{cH(1/c)} + 1 \\ \Leftrightarrow\ & t = \frac{n}{2^{cH(1/c)} + 1} \\ \Leftrightarrow\ & t = \frac{n}{2^{\log(x-1)} + 1}, \text{ by (12)} \\ \Leftrightarrow\ & t = \frac{n}{x}. \end{aligned}$$

Note that $\frac{d^2}{dt^2}M(n,t) = H''(\frac{t}{n})/n < 0$ for all values of t, by *(H3)*. Thus,

$$M(n,t) \leq M\left(n, \frac{n}{x}\right), \text{ for all values of } t\,. \tag{13}$$

The value of $M(n, \frac{n}{x})$ can be calculated as follows:

$$\begin{aligned} M\left(n, \frac{n}{x}\right) &= nH\left(\frac{1}{x}\right) - c\frac{n}{x}H\left(\frac{1}{c}\right) \\ &= n\left(\log(x) + \frac{1-x}{x}\log(x-1) - \frac{c}{x}H(1/c)\right), \text{ by } \textit{(H1)} \\ &= n\left(\log(x) + \frac{1-x}{x}\log(x-1) - \frac{1}{x}\log(x-1)\right), \text{ by (12)} \\ &= n\big(\log(x) - \log(x-1)\big) = n\log\left(\frac{x}{x-1}\right) \\ &= n\log\left(1 + \frac{(c-1)^{c-1}}{c^c}\right). \end{aligned} \tag{14}$$

Combining (11), (13), and (14), we conclude that

$$\log(C(n, \lfloor ct \rfloor, t)) \leq n\log\left(1 + \frac{(c-1)^{c-1}}{c^c}\right) + 3\log(n+1).$$

Lower Bound: By Lemma 1,

$$\log\left(\max_{t\colon \lfloor ct \rfloor < n} C(n, \lfloor ct \rfloor, t)\right) \geq \log\left(\max_{t\colon \lfloor ct \rfloor < n} \frac{\binom{n}{t}}{\binom{\lfloor ct \rfloor}{t}}\right).$$

This proves (5). In order to prove (6), first note that by Lemma 15,

$$\log\left(1 + \frac{(c-1)^{c-1}}{c^c}\right) n \leq \frac{n}{c}\,.$$

Thus, for $c \geq \frac{n}{2}$, the righthand side of (6) is negative, and hence, the inequality is trivially true.

Assume now that $c < \frac{n}{2}$. We will determine an integer value of t such that $\binom{n}{t}/\binom{\lfloor ct \rfloor}{t}$ becomes sufficiently large. First, we use Lemma 16:

$$\frac{\binom{n}{t}}{\binom{\lfloor ct \rfloor}{t}} \geq \frac{2^{nH(t/n)}}{(n+1)\cdot 2^{\lfloor ct \rfloor H(t/\lfloor ct \rfloor)}} = \frac{2^{nH(t/n) - \lfloor ct \rfloor H(t/\lfloor ct \rfloor)}}{n+1}$$

It is possible that $t = \lfloor ct \rfloor$, but this is fine since $H(1) = 0$. Using *(H4)*, we see that

$$\lfloor ct \rfloor H\left(\frac{t}{\lfloor ct \rfloor}\right) \leq ctH\left(\frac{t}{ct}\right) = ctH\left(\frac{1}{c}\right).$$

Thus,

$$\log\left(\frac{\binom{n}{t}}{\binom{\lfloor ct\rfloor}{t}}\right) \geq nH\left(\frac{t}{n}\right) - ctH\left(\frac{1}{c}\right) - \log(n+1) = M(n,t) - \log(n+1). \quad (15)$$

Let $t' = \frac{n}{x}$. We know that $M(n,t)$ attains its maximum value when $t = t'$. Since $c > 1$, it is clear that $x > c$ and hence $t' < \frac{n}{c}$. It follows that $\lfloor ct' \rfloor < n$. However, t' might not be an integer. In what follows, we will first argue that $\lfloor c\lceil t' \rceil \rfloor < n$ and then that $M(n, \lceil t' \rceil)$ is close to $M(n,t')$. The desired lower bound will then follow by setting $t = \lceil t' \rceil$.

Using calculus, it can be verified that, for $c > 1$, x/c is increasing in c. Hence,

$$\begin{aligned}
\frac{x}{c} &= \left(\frac{c}{c-1}\right)^{c-1} + \frac{1}{c} \\
&\geq \lim_{c\to 1^+}\left(\left(\frac{c}{c-1}\right)^{c-1} + \frac{1}{c}\right), \text{ for } c > 1 \\
&= \lim_{c\to 1^+}\left(1 + \frac{1}{c-1}\right)^{c-1} + \lim_{c\to 1^+}\frac{1}{c} = \lim_{a\to 0^+}\left(1+\frac{1}{a}\right)^{a} + 1 = 2.
\end{aligned}$$

Thus, $c \leq x/2$, and hence,

$$\lfloor c\lceil t' \rceil \rfloor \leq c\left\lceil \frac{n}{x} \right\rceil < \frac{cn}{x} + c \leq \frac{n}{2} + c < n.$$

Note that $\frac{d}{dt}M(n,t) < 0$ for $t > t'$, so $M(n, \lceil t' \rceil) \geq M(n, t'+1)$. Combining this observation with (H2) and (H5), we get that

$$\begin{aligned}
M(n, \lceil t' \rceil) \geq M(n, t'+1) &= nH\left(\frac{t'+1}{n}\right) - c(t'+1)H\left(\frac{1}{c}\right) \\
&= nH\left(\frac{1}{x} + \frac{1}{n}\right) - c\frac{n}{x}H\left(\frac{1}{c}\right) - cH\left(\frac{1}{c}\right) \\
&\geq nH\left(\frac{1}{x} + \frac{1}{n}\right) - c\frac{n}{x}H\left(\frac{1}{c}\right) - \log n - 2, \text{ by } \textit{(H2)} \\
&\geq nH\left(\frac{1}{x}\right) - c\frac{n}{x}H\left(\frac{1}{c}\right) - \log n - 5, \text{ by } \textit{(H5)} \\
&= M(n,t') - \log n - 5.
\end{aligned}$$

By choosing $t = \lceil t' \rceil$ in the max, we conclude that

$$\begin{aligned}
\log\left(\max_{t:\, \lfloor ct\rfloor < n} \frac{\binom{n}{t}}{\binom{\lfloor ct\rfloor}{t}}\right) &\geq M(n, \lceil t' \rceil) - \log(n+1), \text{ by (15)} \\
&\geq M(n,t') - \log(n+1) - \log n - 5 \\
&\geq n\log\left(1 + \frac{(c-1)^{c-1}}{c^c}\right) - 2\log(n+1) - 5, \text{ by (14)}.
\end{aligned}$$

□

The following lemma is used for proving Theorem 10.

Lemma 18 *If c is an integer-valued function of n and $c > 1$, it holds that*

$$\log\left(\max_{t\colon ct<n} \frac{\binom{n}{t}}{\binom{ct}{t}}\right) = \Omega\left(\frac{n}{c}\right).$$

Proof Assume that c is an integer-valued function of n, that $c > 1$ and that $ct < n$. It follows that

$$\frac{\binom{n}{t}}{\binom{ct}{t}} = \frac{n!(ct-t)!}{(n-t)!(ct)!} \geq \frac{n(n-1)\cdots(n-t+1)}{(ct)(ct-1)\cdots(ct-t+1)}$$

Let $t = \lfloor \frac{n}{ec} \rfloor$. Then

$$\begin{aligned}\frac{\binom{n}{t}}{\binom{ct}{t}} &= \frac{n(n-1)\cdots(n-t+1)}{(ct)(ct-1)\cdots(ct-t+1)} \geq \frac{n(n-1)\cdots(n-t+1)}{\frac{n}{e}(\frac{n}{e}-1)\cdots(\frac{n}{e}-t+1)} \\ &= \frac{n}{\frac{n}{e}}\frac{n-1}{\frac{n}{e}-1}\cdots\frac{n-t+1}{\frac{n}{e}-t+1} \geq e^t.\end{aligned}$$

Since

$$\log(e^t) = t\log(e) \geq \left(\frac{n}{ec}-1\right)\log e = \frac{n}{e\ln(2)\,c} - \log(e) = \Omega\left(\frac{n}{c}\right),$$

this proves the lemma by choosing $t = \lfloor \frac{n}{ec} \rfloor$. □

A.5 Approximating the Advice Complexity Bounds for MAXASG

Lemma 20 of this section is used for Theorems 6–8. In proving Lemma 20, the following lemma will be useful.

Lemma 19 *For all n, c, it holds that*

$$\max_{u\colon 0<u<n} \frac{\binom{n}{u}}{\binom{n-\lceil u/c\rceil}{n-u}} \leq n\left(\max_{t\colon \lfloor ct\rfloor<n} \frac{\binom{n}{t}}{\binom{\lfloor ct\rfloor}{t}}\right).$$

On the other hand, it also holds that

$$\max_{u\colon 0<u<n} \frac{\binom{n}{u}}{\binom{n-\lceil u/c\rceil}{n-u}} \geq \frac{1}{n}\left(\max_{t\colon \lfloor ct\rfloor<n} \frac{\binom{n}{t}}{\binom{\lfloor ct\rfloor}{t}}\right).$$

Proof Let

$$f_{n,c}(t) = \frac{\binom{n}{t}}{\binom{\lfloor ct\rfloor}{t}} \text{ and } g_{n,c}(u) = \frac{\binom{n}{u}}{\binom{n-\lceil u/c\rceil}{n-u}}.$$

In order to prove the upper bound, we show that $f_{n,c}(\lfloor u/c \rfloor) \geq g_{n,c}(u)/n$, for any integer u, $0 < u < n$. Note that $\lfloor u/c \rfloor < u$, since $c > 1$.

$$\begin{aligned}
f_{n,c}(\lfloor u/c \rfloor) &= \frac{\binom{n}{\lfloor u/c \rfloor}}{\binom{\lfloor c \lfloor u/c \rfloor \rfloor}{\lfloor u/c \rfloor}} \\
&\geq \frac{\binom{n}{\lfloor u/c \rfloor}}{\binom{u}{\lfloor u/c \rfloor}}, \text{ by } (B2) \\
&= \frac{\binom{n}{u}}{\binom{n-\lfloor u/c \rfloor}{n-u}}, \text{ by } (B3) \\
&\geq \frac{u - \lfloor u/c \rfloor}{n - \lfloor u/c \rfloor} \frac{\binom{n}{u}}{\binom{n-\lceil u/c \rceil}{n-u}}, \text{ by } (B1) \\
&\geq \frac{u - \lfloor u/c \rfloor}{n - \lfloor u/c \rfloor} g_{n,c}(u) \\
&\geq \frac{1}{n} g_{n,c}(u), \text{ since } u - \lfloor u/c \rfloor \geq 1.
\end{aligned}$$

By $(B1)$, the second last inequality is actually an equality, unless u/c is an integer.

In order to prove the lower bound, we will show that $g_{n,c}(\lceil ct \rceil) \geq f_{n,c}(t)/n$, for any integer t with $\lfloor ct \rfloor < n$. Note that $t < \lceil ct \rceil$, since $c > 1$.

$$\begin{aligned}
g_{n,c}(\lceil ct \rceil) &= \frac{\binom{n}{\lceil ct \rceil}}{\binom{n-\lceil \lceil ct \rceil / c \rceil}{n-\lceil ct \rceil}} \\
&\geq \frac{\binom{n}{\lceil ct \rceil}}{\binom{n-t}{n-\lceil ct \rceil}}, \text{ by } (B2) \\
&= \frac{\binom{n}{n-\lceil ct \rceil}}{\binom{n-t}{n-\lceil ct \rceil}} \\
&= \frac{\binom{n}{n-t}}{\binom{\lceil ct \rceil}{t}}, \text{ by } (B3) \\
&= \frac{\binom{n}{n-t}}{\frac{\lceil ct \rceil}{\lceil ct \rceil - t}\binom{\lfloor ct \rfloor}{t}}, \text{ by } (B1) \\
&= \frac{\lceil ct \rceil - t}{\lceil ct \rceil} f_{n,c}(t) \\
&\geq \frac{1}{n} f_{n,c}(t), \text{ since } \lceil ct \rceil - t \geq 1 \text{ and } \lceil ct \rceil \leq n.
\end{aligned}$$

□

Lemma 20 *Let $c > 1$ and $n \geq 3$. It holds that*

$$\log\left(\max_{u:\,0<u<n} C\left(n, n-\left\lceil\frac{u}{c}\right\rceil, n-u\right)\right) \geq \log\left(\max_{u:\,0<u<n} \frac{\binom{n}{u}}{\binom{n-\lceil\frac{u}{c}\rceil}{n-u}}\right)$$
$$\geq \log\left(1+\frac{(c-1)^{c-1}}{c^c}\right) n - 3\log n - 6.$$

Furthermore,

$$\log\left(\max_{u:\,0<u<n} C\left(n, n-\left\lceil\frac{u}{c}\right\rceil, n-u\right)\right) \leq \log\left(\max_{u:\,0<u<n} \frac{\binom{n}{u}}{\binom{n-\lceil\frac{u}{c}\rceil}{n-u}} n\right)$$
$$\leq \log\left(1+\frac{(c-1)^{c-1}}{c^c}\right) n + 4\log(n+1)$$

Proof We prove the lower bound first.

$$\log\left(\max_{u:\,0<u<n} C\left(n, n-\left\lceil\frac{u}{c}\right\rceil, n-u\right)\right)$$
$$\geq \log\left(\max_{u:\,0<u<n} \frac{\binom{n}{u}}{\binom{n-\lceil\frac{u}{c}\rceil}{n-u}}\right), \text{ by Lemma 1}$$
$$\geq \log\left(\max_{t:\,\lfloor ct\rfloor<n} \frac{\binom{n}{t}}{\binom{\lfloor ct\rfloor}{t}}\right) - \log n, \text{ by Lemma 19}$$
$$\geq \log\left(1+\frac{(c-1)^{c-1}}{c^c}\right) n - 2\log(n+1) - 5 - \log n, \text{ by (6)}$$
$$\geq \log\left(1+\frac{(c-1)^{c-1}}{c^c}\right) n - 3\log n - 6, \text{ since } n \geq 3.$$

We now prove the upper bound.

$$\log\left(\max_{u:\,0<u<n} C\left(n, n-\left\lceil\frac{u}{c}\right\rceil, n-u\right)\right)$$
$$\leq \log\left(\max_{u:\,0<u<n} \frac{\binom{n}{u}}{\binom{n-\lceil\frac{u}{c}\rceil}{n-u}} \left(1+\ln\binom{n-\lceil u/c\rceil}{n-u}\right)\right), \text{ by Lemma 1}$$
$$\leq \log\left(\max_{u:\,0<u<n} \frac{\binom{n}{u}}{\binom{n-\lceil\frac{u}{c}\rceil}{n-u}} n\right)$$
$$= \log\left(\max_{u:\,0<u<n} \frac{\binom{n}{u}}{\binom{n-\lceil\frac{u}{c}\rceil}{n-u}}\right) + \log n$$
$$\leq \log\left(\max_{t:\,\lfloor ct\rfloor<n} \frac{\binom{n}{t}}{\binom{\lfloor ct\rfloor}{t}} n\right) + \log n, \text{ by Lemma 19}$$
$$\leq \log\left(1+\frac{(c-1)^{c-1}}{c^c}\right) n + 4\log(n+1), \text{ by (8)}$$

□